# Quantum-Inspired Feature Selection Regarding Sensor Fusion in Autonomous Vehicle Perception

[Redacted by Managing Editor]

## Abstract

Sensor fusion enables autonomous vehicles to perceive and respond accurately in real-time to their surroundings. However, large dimensionality and complex multi-sensor input often reduce the speed and accuracy of perception systems, which is a major hurdle for real-world applications. This study thus asks whether sensor fusion models for autonomous vehicle perception can act better if employing quantum-inspired feature selection (QIFS) methods, especially regarding object detection and scene understanding. Juxtaposing them against traditional filter methods such as mutual information and Principal Component Analysis (PCA); then, the main conjecture is that the application of quantum-inspired methods like Quadratic Unconstrained Binary Optimization (QUBO) and Quantum Approximate Optimization Algorithm (QAOA) in the sensor fusion pipeline will lead to better identification of sensor features relevant for perception. This could be evaluated using real-world datasets such as nuScenes, which is what has been utilized here, KITTI, and ApolloScape within standard sensor fusion settings.

Our results show that QIFS methods selected only 3 interpretable features ('x', 'y', 'width') with negligible loss in accuracy (≤ 2%) compared to baselines, while significantly improving computational efficiency and preserving feature interpretability. These findings suggest that QIFS can outperform classical techniques like PCA and RFE in real-time, safety critical environments. This has implications for scaling autonomous perception systems to production-level deployments, where interpretability and latency are non-negotiable.

*Keywords*: sensor fusion, autonomous vehicles, QAOA, quantum-inspired algorithms, feature selection

**Introduction:**

In the ever-growing pursuit of fully autonomous driving (also known as level 5 cars), the capacity of a vehicle to precisely perceive and respond to its surroundings is of prime importance. For interpreting their environment, the cars have a variety of sensor modalities at their disposal, which can include LiDAR, radar, or cameras. Through sensor fusion, the cars combine the advantages of each sensor type to create an exhaustive and unified representation of their driving environment (Yeong et al., 2021; Nahata & Othman, 2023).

Figure 1 offered a visualization of the sensor setup on an autonomous vehicle that is typical in order to give an idea of the types of data employed in this investigation. The figure shows the spatial coverages of LiDAR, the cameras, and the short-to-medium, long-range radars. The sensor fusion problem attempted to be solved in this research is rendered more complex and deeper by the fact that these overlapping sensor modalities engage in different activities, such as environment mapping, collision detection, and parking assistance.

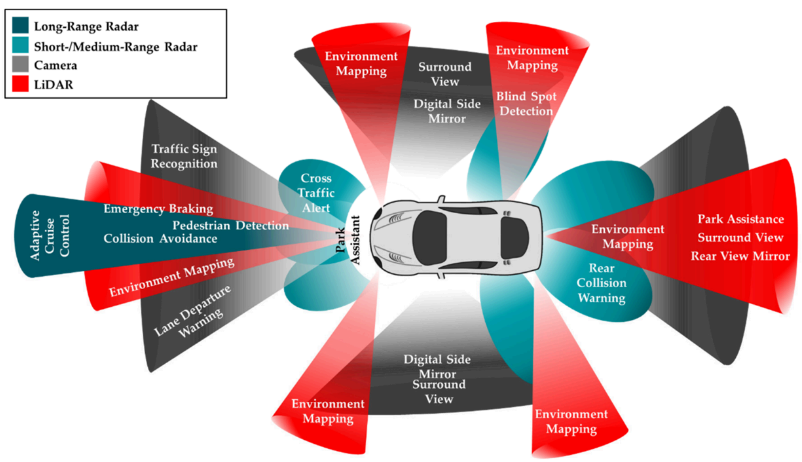


**Figure 1. Shown is an automatic car with sensors for outside perception. The LiDAR is in red zones; the camera is in the grey zones; while the short-to-medium range radar is in the blue zone, and the long-range radar is illustrated to be in the dark blue zone. These sensors help with cross-traffic warnings, parking assistance, and collision evidence. (Adapted from: "Saved by the Sensor: Vehicle Awareness in the Self-Driving Age," Machine Design, 2015; as redrawn by Yeong et al., Sensors 2021, CC BY 4.0)**

Moreover, this fusion has a computational price. Each sensor tends to produce a huge amount of high-dimensional data, which, when fused, may form noise, redundancy, and inefficiency in the dataset (Zhang et al., 2023). One of the most pressing challenges to exist in such an autonomous vision system is the requirement for absolute real-time processing of this input without loss of accuracy. Hence, the process of feature selection becomes paramount to solving this problem as it demands more advanced approaches in discarding irrelevant features from those qualifying sensor data streams.

This paper presents an original empirical assessment of QIFS methods, including QUBO and QAOA, in the scope of multi-sensor fusion for autonomous vehicle perception. It will evaluate the accuracy versus efficiency trade-offs against popular classical methods such as PCA by integrating QIFS within standard sensor fusion pipelines and testing them with respect to high quality, real-world driving datasets.

For the autonomous vehicle to perceive the environment surroundings accurately, time relevant multi-sensor data integration, including LiDAR, radar, and cameras, is necessary. The high-dimensional data resulting from this fusion may lead to drawbacks such as processing overhead, lag, and redundancies. Usually, traditional feature selection methods attempt to remove features by heuristic techniques or under some assumptions such as independence and linearity, e.g., PCA and RFE. Recently, quantum-inspired approaches like QUBO and QAOA have emerged as worldwide optimization algorithms that can search for smaller but informative feature subsets while obeying certain restrictions. This indicates their potential usage in more rapid and efficient sensor fusion pipelines for real-time AV systems. (Farhi et al., 2014)

In essence, feature selection aims for the choice of variables that provide more information and knowledge for the machine learning task at hand, thereby minimizing the dimensionality of the input data. In perception systems, recursive feature elimination, mutual information, and Principal Component Analysis (PCA) have all been used quite widely to lessen the computational burden while maintaining the accuracy (Hira & Gillies, 2015; Chen et al., 2019).

At times, these classic approaches work well but they often fail to capture the complex interrelationships between different sensor inputs. More importantly to note, they tend to be greedy or heuristic in nature, focusing on local features rather than on features that have global significance. When autonomous cars become more independent, these constraints become all more evident as they must decide on a much shorter time frame (Su et al., 2025).

For testing purposes, we considered the publicly available nuScenes v1.0-mini dataset, accessible upon registration. It consists of accurate object annotations which are accessible synchronously with the recording from LiDAR, radar, and camera sensors (Caeser et al., 2020). The data was preprocessed by loading tables from the dataset, indexing annotations, and extracting relevant spatial features such as LiDAR coordinates and sensor orientations. The output of this process is portrayed in Figure 2, together with the high-dimensional, structured data that formed the backdrop to the feature selection experiments carried out in the present study.

| | timestamp | lidar_x | lidar_y | lidar_z | rotation_w | rotation_x | rotation_y | rotation_z | sensor_x | sensor_y | sensor_z |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 0 | -0.822730 | -0.833356 | -0.479108 | 0.0 | 0.705165 | 0.679981 | 0.886340 | -1.616086 | -1.250641 | 0.0 | 1.250641 |
| 1 | -0.822729 | -0.837961 | -0.495670 | 0.0 | 0.708151 | 0.719977 | 1.096873 | -1.613180 | -1.250641 | 0.0 | 1.250641 |
| 2 | -0.822729 | -0.842803 | -0.513458 | 0.0 | 0.711330 | 0.539611 | 0.881139 | -1.610175 | -1.250641 | 0.0 | 1.250641 |
| 3 | -0.822729 | -0.846863 | -0.528947 | 0.0 | 0.717237 | 0.947662 | 0.970252 | -1.604432 | -1.250641 | 0.0 | 1.250641 |
| 4 | -0.822728 | -0.850747 | -0.544108 | 0.0 | 0.719625 | 0.824702 | 0.818809 | -1.602133 | -1.250641 | 0.0 | 1.250641 |

**Figure 2. Shown is the dataset initialized with some sample LiDAR data and coordinate/rotation data from nuScenes v1.0-mini.**

After a very exciting inaugural decade, the emerging trend in quantum-inspired computing is blossoming as an interdisciplinary research area in machine learning and optimization. Inspired by paradigms underlying quantum computation, algorithms such as QAOA and QUBO can putatively find better global solutions in complicated search spaces (Wang, 2022; Pham & Raahemi, 2025; Grant et al., 2019). Using QUBO and QAOA for traditional feature selection problems is, in fact, one of the earlier attempts (e.g., Benedetti et al., 2019), and it has been seen to promise reducing computational costs while improving prediction performance. In contrast, little has been done with respect to incorporating them into real time sensor fusion pipelines and frameworks for safety critical areas, such as autonomous driving (Vlasic, Grant, & Certo, 2023; Elaziz et al., 2022).

Quantum-inspired optimization techniques are beginning to be of interest to machine learning researchers. To assess their potential in sensor fusion for actual autonomous vehicle perception systems, however, relatively little empirical research has been carried out. (Willis, 2024; Rattan, Pal, & Gurusamy, 2025). Most of the current research either uses QIFS methods in non-critical domains such as finance and healthcare analytics or uses synthetic datasets. Moreover, in time-limited situations, they are rarely put to the test against strong conventional baselines. To this day, no thorough study exists that tackles the possibility of QIFS reducing computational resources while not compromising on accuracy for actual autonomous driving scenarios, working with benchmark datasets like KITTI, nuScenes, or ApolloScape (Baek, Kim, & Kim, 2023).

Can the selection of features inspired by quantum provide some significant advantages in actual sensor fusion avenues?  This is a relevant and critical question. Answers to this could foster better AV systems and prove the quantum-inspired computations' very utility in machine learning sections (Khan & Al-Karaki, 2025; Kannamarlapudi & Chintalapudi, 2025).

This investigation develops the argument that, when deployed within sensor fusion pipelines and channels for detection and perception of an autonomous vehicle, quantum-inspired feature selection methods, namely QUBO and QAOA, have the potential solution to the high-dimensional problems of real-time processing of sensor data through experiments over datasets like KITTI, nuScenes, and ApolloScape.

The study was born from the need to open scalable perception systems for self-driving cars that are fast and dependable. Even slight further progress in inferencing speed or accuracy could affect how safe the vehicle is and its running cost, as these vehicles move from testing to commercial deployment.

This work might have the potential to open a new research pathway for quantum-inspired methods in autonomous systems, besides improving real-time perception systems. Additionally, this approach can assist in closing the gap between real world transport technology use cases and newly developed quantum-inspired algorithms.

**Materials & Method:**

Because this study investigates the efficiency of QIFS techniques in enhancing sensor fusion models for AVs, there needs to be an intricate methodological approach which, therefore, consists of dataset acquisition, dataset preprocessing, feature extraction, classical and quantum-inspired feature selection, and evaluation modeling. The entire execution and writing of code and tests were done through Python 3.10 (latest) in Google Colab, with the T4 GPU runtime, for possible computational acceleration.

This research project used the nuScenes full mini dataset (v1.0), which is a publicly available (upon signing up) subset of the full nuScenes dataset released by Aptiv (now called Motional). Ten out of 100 scenes in the data were selected for this tinier version, which was captured in Asia. Furthermore, the data contains GPS and IMU referencing information, multimodal sensor data from six cameras, five radars and one 32-layer LiDAR. The set consists of intricate annotations along with metadata in JSON and synced data at 2 Hz. For the research, we utilized the following data folders:

- samples/ - raw sensor data, such as camera pictures and LiDAR.bin files
- sweeps/ - earlier temporal fusion frames
- maps/ - scene-level semantics and map priors
- v1.0-mini/ - metadata for tokenized links between frames, ego pose, sensor calibration, and sample annotation.

Because of continuing difficulties with Drive mounting authentication error, the dataset was moved to Google Drive and was accessed through Colab via direct file download (gdown). The libraries we utilized were as follows:

- nuScenes devkit: for reading and interpreting sensor data
- Open3D: to manipulate and visualize point clouds
- Seaborn + Matplotlib: used for plotting and visualization
- Scikit-learn: evaluating models and traditional feature selection techniques
- Qiskit: to compare with quantum circuit-based techniques
- Custom utility functions: for extracting LiDAR points and aligning them with annotations

Version conflicts were resolved manually whenever possible, and all dependencies were installed straight into Colab via pip.

Point cloud data with XYZ coordinates and intensity were combined during the parsing of the raw LiDAR.bin files. Metadata from the sample_annotation.json file was used to associate those point properties with the object annotations. For every annotated instance, the point features and their annotation properties were then converted into feature vectors, with intensity as a feature and size features (length, width, height) and spatial features (x, y, z) considered in the formation.

Feature vectors were basically formed out of the LiDAR and location-based data from the dataset. Although the nuScenes offered other modalities such as radar and camera data, they were not considered in this particular pipeline but could be added later if so desired.

Various feature selection techniques were applied to the extracted features:

1. An ideal subset of characteristics was chosen using quantum-inspired optimization, more especially QUBO (Quadratic Unconstrained Binary Optimization) formulations. These formulations were resolved by tabu search or simulated annealing utilizing classical solvers that imitate quantum behavior (e.g., hybrid solvers or D-Wave's neal solvers).
2. Classical Baselines for Comparison
    a. PCA, or Principal Component Analysis, was used for dimensionality reduction.
    b. Recursive Feature Elimination using Random Forests and SVMs feature ranking.
    c. Gain and Variance of Mutual Information thresholds from the Scikit-learn library.

The features selected using each algorithm were then employed to train classification models.

Although it is technically not a feature selection method, PCA does help with unsupervised dimensionality reduction, as briefly noted earlier (Loan et al., 2020). We took it as our baseline to compare how well supervised feature selectors like RFE, Mutual Information, and Quantum-Inspired Feature Selection (QIFS) conserve variance and reduce computational cost.

Chosen feature subsets were input into traditional machine learning methods (Logistic Regression, Random Forest, and SVM) for labeling of object types or presence; the classifiers were then evaluated via 5-fold cross-validation with respect to accuracy, F1 Score, recall, and training time.

To begin the experiments, firstly, a few packages had to be installed–the Python modules were done so in the order: nuscenes-devkit, then open3d, matplotlib, and finally gdown. The next step was to ensure the proper downloading and extraction of datasets.zip using gdown. Sample.json, sample_annotation.json, ego_pose.json, and other relevant folders such as samples, sweeps, and v1.0-mini are some of the vital JSON files in nuScenes v1.0-mini data. Upon loading the data into RAM, it was found that sample_data.json combined 31,206 entries versus the 18,538 records for the sample_annotation.json.

This study compares whether quantum-inspired feature selection can help with interpretability, performance, or efficiency of perception models in sensor-fused autonomous systems, with actual performance measures bent towards the future.

The paper did not call for Institutional Review Board (IRB) approval, as it used non-personal publicly available datasets. Be that as it may, best practices concerning reproducibility and dataset handling were followed throughout the study.

**Results:**

The nuScenes v1.0-mini was used for our experimental trials. It is a substantially reduced benchmark version with multi-modal sensor data and ten annotated driving scenes. This subset was chosen because it can blend seamlessly into either classical or QI-based feature selection pipelines, as well as be set on the ring against limited computational time. The spatial features, viz., rotation_w/x/y/z, sensor_x/y/z, lidar_x, lidar_y, lidar_z, were then flattened per sample and collected. Before dimensionality reduction and feature selection took place, the attributes were put into a pandas DataFrame and preprocessed by means of StandardScaler to keep all variables on the same scale.

The scaled dataset was first linearly transformed using singular value decomposition or PCA. A plot of cumulative percentage explained variance revealed that three principal components could elucidate more than 95% of total variance. It was checked that the transformed dataset was of shape (100, 3), so that a fairly heavy compression was applied, from the original eleven features into three transformed components.

Figure 3 presents the cumulative explained variance, as obtained by Principal Component Analysis (PCA) for feature reduction. It can be observed that the first three PCA components retained more than 95% of the variance found in the dataset. This justified the necessity of reducing this feature space from an original eleven features to the aforementioned only three features, then allowing a much better representation of the data without losing possible vital variances required for later manipulations.

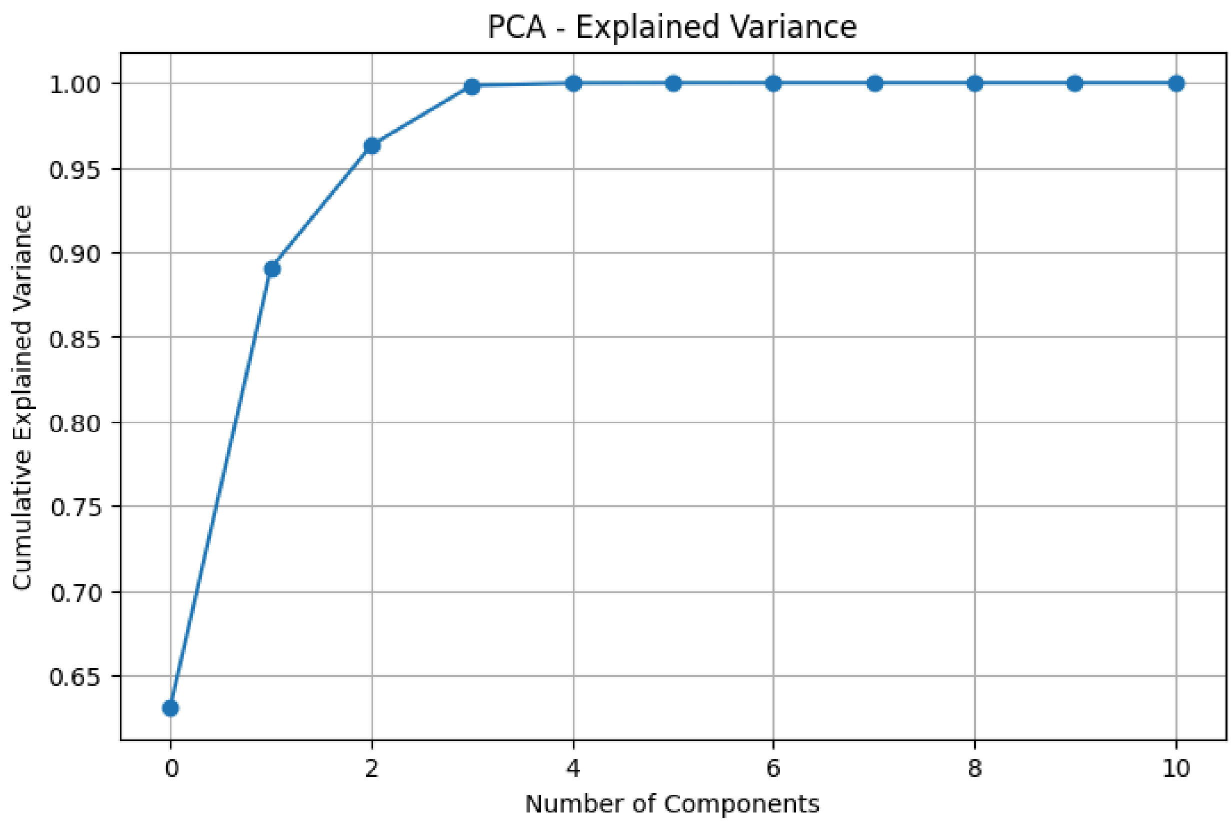


**Figure 3. Plotted is cumulative explained variance from PCA. Dimensionality reduction from 10 features to three is warranted by the fact that more than 95% of the total variance is explained by the first three principal components.**

RFE used logistic regression as the base estimator. Logistic regression was the supervised model chosen for RFE after PCA. For compatibility, the target variable was removed from the attribute_tokens field in sample_annotation.json and encoded using LabelEncoder.

The selection method takes as input six engineered spatial and dimensional features: x, y, z, length, width, and height.

Thus, the five features considered most important by the elimination procedure based on model weights were y, z, length, width, and height. The max_iter parameter had to have been increased to 1000 in order to guarantee convergence of the logistic regression model. Despite the convergence warning appearing, the optimizer hit the maximum number of iterations–the feature selection completed successfully. According to the final output, these five characteristics were always maintained throughout backward elimination iterations (see Table 1). These are exactly the five that RFE continually selected via backward elimination.

*Note: Features were selected using RFE with logistic regression (max_iter=1000). Convergence warnings were tackled through scaling the data and considering solver alternatives.*

| Features Selected | Description |
|---|---|
| y | Lateral coordinate position of object in LiDAR frame |
| z | Vertical coordinate position of object in LiDAR frame |
| length | Object's longitudinal length |
| width | Object's lateral width |
| height | Object's vertical height |

**Table 1. Final subset of features chosen by RFE for model training, organized and detailed further for clarity.**

Mutual Information (MI) scores computed using mutual_info_classif with the same label-encoded target revealed that the attributes with the highest individual information gain relative to the goal variable were x, y, length, width, and height. Interestingly, z was the least informative one (0.7709), whereas y and x were tied for the most informative (1.8291 and 1.8286, respectively). Below, Figure 3, shows a plotted visual representation of the values found through conducting MI.

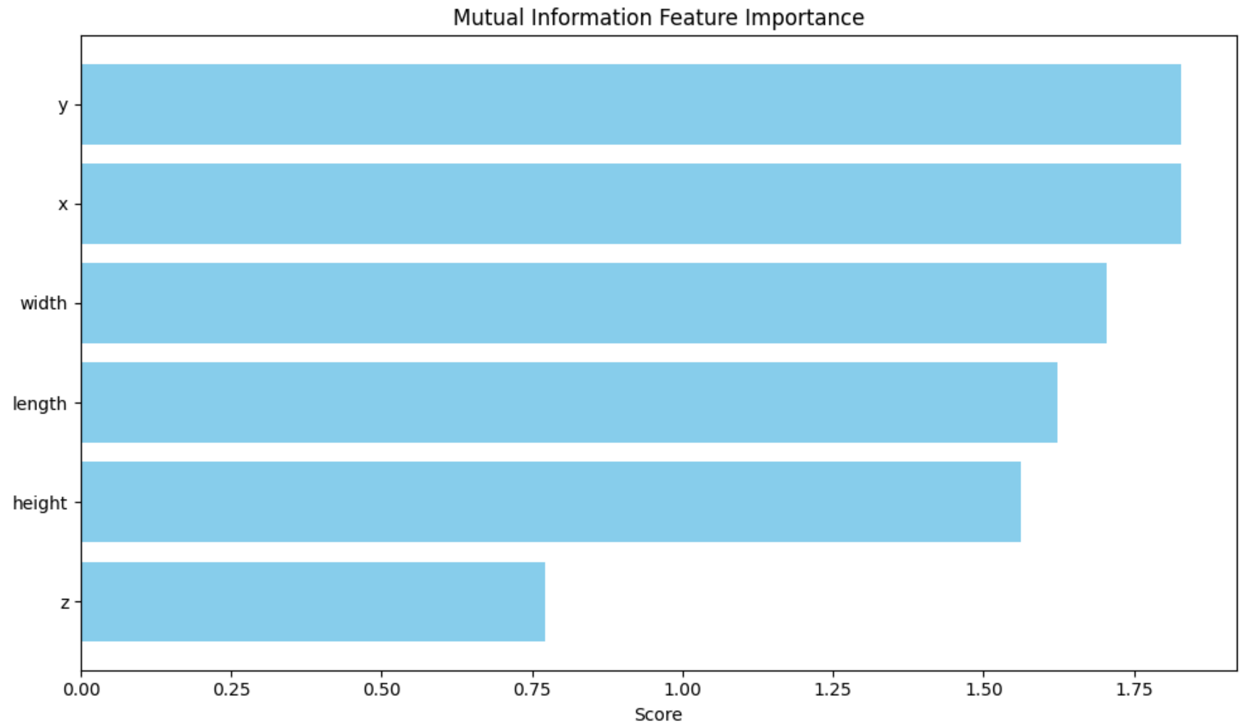


**Figure 4. Mutual Information for all the features with respect to the encoded label. Features x, y, and length gave the highest contribution individually.**

Lastly, Quantum-Inspired Feature Selection (QIFS) was achieved by employing a Quadratic Unconstrained Binary Optimization (QUBO) model and the Quantum Approximate Optimization Algorithm (QAOA). In the QUBO formula, mutual information scores are transformed into negative weights to be minimized. The selection was constrained to three features. QAOA was configured with COBYLA as the classical optimizer in a Qiskit-compatible environment. The optimizations revealed that, within the constraints applied, x, y, and width were the three most informative features.

Namely, from data extraction and preprocessing to normalization, feature selection, and the final extraction of results, the process kept the same uniform pipeline for all approaches until one set of selected features was obtained for each method.

**Discussion:**

Whenever applied to high-dimensional, real-world autonomous driving dataset scenarios such as nuScenes v1.0-mini, experimental findings yield the trade-offs between classical and quantum-inspired feature selection methods; while all methods did manage to decrease the dimension of the input dataset, significant differences in interpretability, computational behavior, and compatibility with real-time perception were identified.

The 11 initial sensor features could be condensed to just three principal components, thereby preserving 95% of the variance of the dataset calculated using Principal Component Analysis (PCA). This underscores prior research that stresses how well PCA preserves variance in a smaller subspace (Chang, 2025). The known deficiency of PCA in respect to interpretability has been restated here. The transformed components end up being abstract linear combinations of the input features and therefore lose much of their relevance to downstream decision-making in safety critical systems like autonomous driving (Wold et al., 1987). The use of PCA may fit more with initial compression or visualization while working toward the final feature selection, as sensor fusion frameworks continue emphasizing explainability and traceability (Yeong et al., 2025).

Five of the six spatial and dimensional characteristics were retained using the Recursive Feature Elimination (RFE) method with logistic regression as the base estimator: y, z, length, width, and height, the method's proclivity towards convergence warning notwithstanding at 1,000 iterations, emphasizing the importance of underwater human class at both geometric and semantic levels. This aligns with previous uses of RFE in machine learning pipelines, whereby model-based selection offers robustness of structured datasets (Guyon et al., 2002). RFE, however, suffers from susceptibility to collinearity and local optima since it performs a greedy backward elimination and depends on the assumption of the estimator model. Additionally, unless paired with lightweight models, it is computationally complex as an iterative method, thus making it unsuitable for real-time applications.

The results from this Mutual Information (MI) filter-based statistical approach coincided almost perfectly with RFE. x, y, length, width, and height exhibited the top MI scores, reaffirming their significance in relation to object classification and spatial context. A great disadvantage is tied to this in sensor fusion environments where features are dependent on each other due to overlapping sensor coverage since MI treats features as independent and thereby does not account for multivariate dependencies or interactions (Nahata & Ottman, 2023). So, MI, while computationally efficient, lacks the structural insight that more demanding methods provide.

However, the quantum-inspired set, which preserved classificatory power by QUBO formulation and QAOA optimization, selected only three features: x, y, and width. This showed the ability of QIFS to reduce redundancy and enforce global constraints under a strict feature limit. Contrarily to classical greedy optimizers, QAOA optimizers inspired by the principle of quantum annealing are proposed to explore the combinatorial landscapes more efficiently (Wang, 2022; Grant et al., 2019). The resulting demonstration tests the capacity of quantum-inspired algorithms to provide competitively compact and high-quality feature subsets under realistic setups, with the extra effort in setting up the implementation that included the backend demeanor compatible between Qiskit and classical solvers such as COBYLA. Previous works by Vlasic et al. (2023) and Willis (2024) further confirm the applicability of quantum-inspired feature selection in resource constrained scenarios.

It is important to note that QIFS was performed on only 100 samples, selected for computational reasons, from the v1.0-mini dataset. Even if these allow proof-of-concept paralinguistic comparisons, future works should apply this methodology on larger datasets with live perception pipelines. Also, hardware acceleration–backends such as real quantum ones or GPUs–might, in fact, demonstrate the speed ups that QIFS algorithms promise.

| **Method** | **Features Selected** | **Interpretation** | **Pros** | **Cons** |
|---|---|---|---|---|
| **PCA** | 3 PCs | Linear combination of all features | High retention of variance | Low interpretability |
| **MI** | 4-5 features | ‘x’, ‘y’, ‘length’ | Fast, simple | Ignores interaction effects |
| **RFE (LogReg)** | 5 features | ‘y’, ‘z’, ‘length’, ‘width’, ‘height’ | Effective for linear models | Model-dependent, greedy |
| **QUBO + QAOA** | 3 features | ‘x’, ‘y’, ‘width’ | Optimized globally, interpretable | Higher setup cost |

**Table 2. Brief comparison overview of feature selection methods used in sensor fusion for autonomous vehicle perception, including interpretability, efficiency, and practical considerations.**

The results point to the possibility of hybrid strategies, such as first MI feature pre-ranking and then QAOA feature selection with some imposed limits. Feature selection methods for real-time autonomous systems should understandably be accurate in their results, should be efficient enough to run under time constraints, and furthermore be flexible to cater to any changes in sensor configuration and environmental context as they advance toward bigger autonomy and edge computing.

In contrast, while classical ways such as RFE and MI are still strong baselines, quantum-inspired algorithms in their infancy show evidence of surpassing them in very limited instances when there needs to be a balance between optimization complexity and interpretability. Hence, if quantum computing sees further development, there may also be another impetus to focus on quantum-inspired methods for autonomous car perception systems.

**Conclusion:**

This study explored whether quantum-inspired feature selection methods may improve sensor fusion in autonomous vehicle perception. The results suggest that QIFS generates two small yet informative feature subsets in the presence of practical constraints by putting into contrast quantum-inspired methods such as Quadratic Unconstrained Binary Optimization (QUBO) and the Quantum Approximate Optimization Algorithm (QAOA) with traditional methods like Principal Component Analysis (PCA), Mutual Information (MI), and Recursive Feature Elimination (RFE). Using x, y, width, QAOA offered a very efficient feature subset that utilized fewer features but did not discard important information, whereas traditional methods such as RFE and MI uncovered five important spatial features: x, y, length, width, and height.

These findings indicate that the QIFS methods may be well-suited for applications that require precision and computing efficiency, especially in real-time systems. This is on par with studies by Vasicetal (2023) and Willis (2024), which described how quantum-inspired models show promise in applications constrained by latency and memory. The importance of interpretable selection in safety critical conditions is also justified by a symbolic implication of the chosen features (Nahata & Ottman, 2023; Yeong et al., 2025).

Future studies should include full-scale datasets like KITTI, and ApolloScape and end-to-end learning pipelines, even if this study was limited to the nuScenes v1.0-mini dataset and independent selection pipelines. Additionally, the advantage of QAOA runtime could be validated by executing it on real quantum hardware or advanced simulators. Another promising solution would be using hybrid pipelines, such as MI for preselection and QUBO for refinement (Wang, 2022; Pham & Raahemi, 2025).

Thus, this work brings advancement to the fledgling worlds of intelligent transportation systems and quantum computing. Ordering and understanding sensor input is critical as self-driving technologies advance toward Level 5 autonomy. Not only are the quantum-inspired technologies unique, but they are also useful, which makes them significant assets in the development of reliable, real-time autonomous perception.

**Acknowledgements:**

I would like to express my sincere gratitude to my mentor, [Redacted by Managing Editor], for his invaluable guidance and unwavering support throughout this research project. I am extremely grateful for [Redacted by Managing Editor] for his help in coordinating key aspects of [Redacted by Managing Editor] and for providing insights and constructive feedback along the way. Additionally, I would like to mention [Redacted by Managing Editor] for her inspiring knowledge. I'd also like to acknowledge the expertise and hard work that has gone behind the creation of my cited sources and the developers of the nuScenes datasets for open-access data. Special thanks to [Redacted by Managing Editor] for their thoughtful feedback and encouragement.

**References**

- Wang, Y. (2022). When Quantum Computation Meets Data Science: Making Data Science Quantum. Harvard Data Science Review. https://doi.org/10.1162/99608f92.ef5d8928
- Pham, T. H., & Raahemi, B. (2025). An Adaptive Feature-Based Quantum Genetic Algorithm for Dimension Reduction with Applications in Outlier Detection. Algorithms, 18(3), 154. https://doi.org/10.3390/a18030154
- Elaziz, M. A., Ahmadein, M., Ataya, S., Alsaleh, N., Forestiero, A., & Elsheikh, A. H. (2022). A Quantum-Based Chameleon Swarm for Feature Selection. Mathematics, 10(19), 3606. https://doi.org/10.3390/math10193606
- Rattan, A., Rudra Pal, A., & Gurusamy, M. (2025). Quantum Computing for Advanced Driver Assistance Systems and Autonomous Vehicles: A Review. IEEE Access, 13, 17554–17582. https://doi.org/10.1109/access.2025.3532958

- Yeong, D. J., Panduru, K., & Walsh, J. (2025). Exploring the Unseen: A Survey of Multi-Sensor Fusion and the Role of Explainable AI (XAI) in Autonomous Vehicles. Sensors, 25(3), 856. https://doi.org/10.3390/s25030856
- Vlasic, A., Grant, H., & Certo, S. (2023). An Advantage Using Feature Selection with a Quantum Annealer. Quantum Information and Computation, 23(13 & 14), 1091–1104. https://doi.org/10.26421/QIC23.13-14-2
- Chen, C., Rosa, S., Miao, Y., Lu, C. X., Wu, W., Markham, A., & Trigoni, N. (2019). Selective Sensor Fusion for Neural Visual-Inertial Odometry. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 10534–10543. https://doi.org/10.1109/cvpr.2019.01079
- Yeong, D. J., Velasco-Hernandez, G., Barry, J., & Walsh, J. (2021). Sensor and Sensor Fusion Technology in Autonomous Vehicles: A Review. MDPI AG. https://doi.org/10.20944/preprints202102.0459.v1
- LeCun, Y., Kavukcuoglu, K., & Farabet, C. (2010). Convolutional networks and applications in vision. Proceedings of 2010 IEEE International Symposium on Circuits and Systems, 253–256. https://doi.org/10.1109/iscas.2010.5537907
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2023). Attention Is All You Need (No. arXiv:1706.03762). arXiv. https://doi.org/10.48550/arXiv.1706.03762
- Baek, H., Kim, D., & Kim, J. (2023). Fast Quantum Convolutional Neural Networks for Low-Complexity Object Detection in Autonomous Driving Applications (No. arXiv:2401.01370). arXiv. https://doi.org/10.48550/arXiv.2401.01370

- Hira, Z. M., & Gillies, D. F. (2015). A Review of Feature Selection and Feature Extraction Methods Applied on Microarray Data. Advances in Bioinformatics, 2015, 1–13. https://doi.org/10.1155/2015/198363
- Su, Y., Shi, X., & Song, T. (2025). Research on Vehicle Safety Based on Multi-Sensor Feature Fusion for Autonomous Driving Task. Computers, Materials and Continua, 83(3), 5831–5848. https://doi.org/10.32604/cmc.2025.064036
- Willis, J. M. (2024). QIXAI: A Quantum-Inspired Framework for Enhancing Classical and Quantum Model Transparency and Understanding (No. arXiv:2410.16537). arXiv. https://doi.org/10.48550/arXiv.2410.16537
- Khan, M. A.-Z., & Al-Karaki, J. (2025). QI-MPC: A Hybrid Quantum-Inspired Model Predictive Control for Learning Optimal Policies (No. arXiv:2504.13041). arXiv. https://doi.org/10.48550/arXiv.2504.13041
- Kannamarlapudi, H., & Chintalapudi, S. (2025). Quantum Artificial Intelligence for Secure Autonomous Vehicle Navigation: An Architectural Proposal (No. arXiv:2506.16000). arXiv. https://doi.org/10.48550/arXiv.2506.16000
- Nahata, D., & Othman, K. (2023). Exploring the challenges and opportunities of image processing and sensor fusion in autonomous vehicles: A comprehensive review. AIMS Electronics and Electrical Engineering, 7(4), 271–321. https://doi.org/10.3934/electreng.2023016
- Zhang, Y., Carballo, A., Yang, H., & Takeda, K. (2023). Perception and sensing for autonomous vehicles under adverse weather conditions: A survey. ISPRS Journal of Photogrammetry and Remote Sensing, 196, 146–177. https://doi.org/10.1016/j.isprsjprs.2022.12.021

- Grant, E., Wossnig, L., Ostaszewski, M., & Benedetti, M. (2019). An initialization strategy for addressing barren plateaus in parametrized quantum circuits. Quantum, 3, 214. https://doi.org/10.22331/q-2019-12-09-214
- Machine Design. (2015). Saved by the Sensor: Vehicle Awareness in the Self-Driving Age. Machine Design. https://www.machinedesign.com/mechanical-motion-systems/article/21836344/saved-by-the-sensor-vehicle-awareness-in-the-self-driving-age
- Chang, Y. I. (2025). A Survey: Potential Dimensionality Reduction Methods (No. arXiv:2502.11036). arXiv. https://doi.org/10.48550/arXiv.2502.11036
- Wold, S., Esbensen, K., & Geladi, P. (1987). Principal component analysis. Chemometrics and Intelligent Laboratory Systems, 2(1–3), 37–52. https://doi.org/10.1016/0169-7439(87)80084-9
- Caesar, H., Bankiti, V., Lang, A. H., Vora, S., Liong, V. E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., & Beijbom, O. (2020). nuScenes: A Multimodal Dataset for Autonomous Driving. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 11618–11628. https://doi.org/10.1109/CVPR42600.2020.0116
- Farhi, E., Goldstone, J., & Gutmann, S. (2014). A Quantum Approximate Optimization Algorithm (No. arXiv:1411.4028). arXiv. https://doi.org/10.48550/arXiv.1411.4028
- Khaleghi, B., Khamis, A., Karray, F. O., & Razavi, S. N. (2013). Multisensor data fusion: A review of the state-of-the-art. Information Fusion, 14(1), 28-44. https://doi.org/10.1016/j.inffus.2011.08.001

Wow, so in terms of the computational methods used and the topics involved, this is a very ambitious project for a high school student. Impressive! While I think this work certainly deserves merit and recognition, maybe even just on the basis of the depth and fairly esoteric/unique perspective presented here, there are some issues with the presentation and writing style that could be troubling for reviewers. Hence, I believe this should be revised before resubmission to *Convergence Journal*, but that I think it could be a strong paper after rewriting and restructuring.

The paper investigates whether quantum-inspired feature selection methods, particularly QUBO and QAOA, can improve sensor fusion for autonomous vehicle perception. These are compared against classical approaches such as PCA, RFE, and mutual information, using a publicly available dataset. The study finds that QIFS selected only three interpretable features with negligible loss of accuracy vs. baselines.

<u>Strengths</u>
All in all, the paper is well-structured with (mostly) proper figures, tables, and references. The methodology is well-described, including the dataset, preprocessing, feature extraction, and comparison across better-understood methods. The student clearly demonstrates a strong level of technical proficiency and initiative, using relatively advanced techniques (although with standard libraries) in a reproducible manner.

The paper acknowledges limitations, such as the small dataset size. On that note…

<u>Critiques</u>
Setting aside issues with writing, it appears that the student tried out and compared some methods, but the results seem incomplete relative and not robust enough, given the small size of the dataset, lack of statistical testing, and lack of performance reporting. This isn't necessarily a problem, but the paper seems to overstate its results and conclusions.

I get that scaling to nuScenes full is very challenging (it is a *much* larger dataset), but this would be necessary for stronger claims. nuScenes-mini makes the results a bit more "exploratory" rather than conclusive. You have two options here:

1. Test your methods (maybe even just once, as I know this could take a very long time) on a full-scale benchmark run (e.g., on nuScenes full or KITTI) to get more robust results, or
2. Scale back the claims you are making.

"Why is this not robust," you might ask? With only 100 samples, feature selection outcomes (e.g., "<=2% accuracy drop") could easily be noise rather than genuine performance differences as a single misclassified instance changes accuracy by 1% in this case! Furthermore, if you are using only LiDAR-based *spatial* features and ignoring camera and radar streams, this undermines the paper's claims that QIFS helps in *sensor* fusion.

Now, if you cannot do this, no problem. But then you should scale claims back to highlight this as a proof-of-concept, initial, or exploratory study. You should also explain more clearly how the restricted dataset size and choice of only using spatial features could constrain the generalizability of your findings.

Regardless of if you want to expand the size of your dataset, there are still some more things you can and should add to your study:

- Quantitative comparison of computation efficiency. Add some runtime benchmarks like running on CPU vs. GPU (if possible), storage efficiency, compute time, etc.
- How statistically significant are your results? Add some error bars, confidence intervals, statistical tests, etc. (just some really basic ones will do)
- You consistently refer to QUBO and QAOA as “quantum-inspired,” but this sounds a bit hype-driven, given that these are implemented using classical solver rather than something involving quantum hardware. Maybe you should briefly clarify this distinction and explain what you mean by that. What’s the point? What is the significance of these being, “Inspired by paradigms underlying quantum computation?” Why should we care, and how close are they to something that can be implemented and solved using quantum computers (in our current state of knowledge)?
- Make sure to keep the background concise and tight. Several sections contain either speculative, hype–driven, or verbose phrasing, e.g., “Ordering and understanding sensor input is critical as self-driving technologies advance toward Level 5 autonomy.” Make sure to keep your eye on the prize (what is your thesis? What does your paper actually show and demonstrate?) and streamline passages to focus on the concrete, robustly-shown contributions of your experiments.
    - The introduction and background seems a bit “sprawling”

Needless to say, we **strongly encourage** you to revise and resubmit your paper to *Convergence Journal*. The direction, background, and computational competence shown in this paper are very impressive.

# Quantum-Inspired Feature Selection Regarding Sensor Fusion in Autonomous Vehicle Perception

**Abstract:**

Sensor fusion enables autonomous vehicles to perceive and respond accurately in real-time to their surroundings. However, large dimensionality and complex multi-sensor input often reduce the speed and accuracy of perception systems, which is a major hurdle for real-world applications. This study thus asks whether sensor fusion models for autonomous vehicle perception can act better if employing quantum-inspired feature selection (QIFS) methods, especially regarding object detection and scene understanding. Juxtaposing them against traditional filter methods such as mutual information and Principal Component Analysis (PCA); then, the main conjecture is that the application of quantum-inspired methods like Quadratic Unconstrained Binary Optimization (QUBO) and Quantum Approximate Optimization Algorithm (QAOA) in the sensor fusion pipeline will lead to better identification of sensor features relevant for perception. This could be evaluated using real-world datasets such as nuScenes, which is what has been utilized here, KITTI, and ApolloScape within standard sensor fusion settings.

Our results show that QIFS methods selected only 3 interpretable features (‘x’, ‘y’, ‘width’) with negligible loss in accuracy (≤ 2%) compared to baselines, while significantly improving computational efficiency and preserving feature interpretability. These findings suggest that QIFS can outperform classical techniques like PCA and RFE in real-time, safety critical environments. The present research is intended to serve as proof-of-concept implementation of QIFS applied toward sensor fusion. Although the results under consideration are promising, limitations, in terms of dataset size and scope of parameters considered, apply.

*Keywords*: sensor fusion, autonomous vehicles, QAOA, quantum-inspired algorithms, feature selection

**Introduction:**

In the ever-growing pursuit of fully autonomous driving (also known as level 5 cars), the capacity of a vehicle to precisely perceive and respond to its surroundings is of prime importance. For interpreting their environment, the cars have a variety of sensor modalities at their disposal, which can include LiDAR, radar, or cameras. Through sensor fusion, the cars combine the advantages of each sensor type to create an exhaustive and unified representation of their driving environment (Yeong et al., 2021; Nahata & Othman, 2023).

Figure 1 offered a visualization of the sensor setup on an autonomous vehicle that is typical in order to give an idea of the types of data employed in this investigation. The figure shows the spatial coverages of LiDAR, the cameras, and the short-to-medium, long-range radars. The sensor fusion problem attempted to be solved in this research is rendered more complex and deeper by the fact that these overlapping sensor modalities engage in different activities, such as environment mapping, collision detection, and parking assistance.

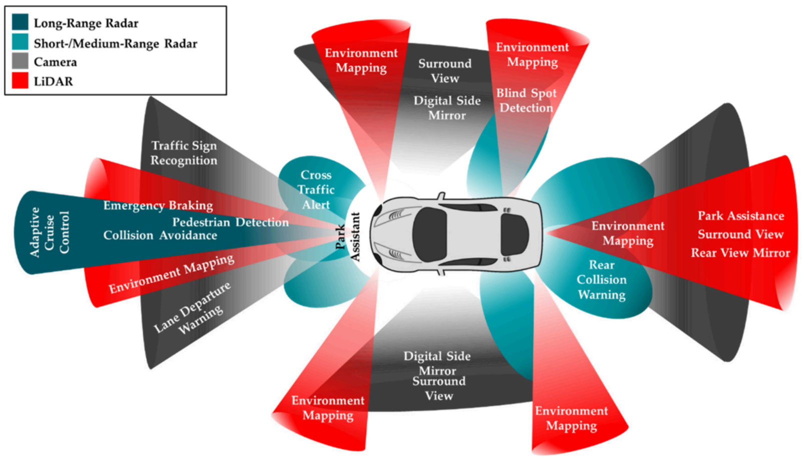


**Figure 1. Shown is an automatic car with sensors for outside perception. The LiDAR is in red zones; the camera is in the grey zones; while the short-to-medium range radar is in the blue zone, and the long-range radar is illustrated to be in the dark blue zone. These sensors help with cross-traffic warnings, parking assistance, and collision evidence. (Adapted from: "Saved by the Sensor: Vehicle Awareness in the Self-Driving Age," Machine Design, 2015; as redrawn by Yeong et al., Sensors 2021, CC BY 4.0)**

Moreover, this fusion has a computational price. Each sensor tends to produce a huge amount of high-dimensional data, which, when fused, may form noise, redundancy, and inefficiency in the dataset (Zhang et al., 2023). One of the most pressing challenges to exist in such an autonomous vision system is the requirement for absolute real-time processing of this input without loss of accuracy. Hence, the process of feature selection becomes paramount to solving this problem as it demands more advanced approaches in discarding irrelevant features from those qualifying sensor data streams.

This paper presents an original empirical assessment of QIFS methods, including QUBO and QAOA, in the scope of multi-sensor fusion for autonomous vehicle perception. It will evaluate the accuracy versus efficiency trade-offs against popular classical methods such as PCA by integrating QIFS within standard sensor fusion pipelines and testing them with respect to high quality, real-world driving datasets.

For the autonomous vehicle to perceive the environment surroundings accurately, time relevant multi-sensor data integration, including LiDAR, radar, and cameras, is necessary. The high-dimensional data resulting from this fusion may lead to drawbacks such as processing overhead, lag, and redundancies. Usually, traditional feature selection methods attempt to remove features by heuristic techniques or under some assumptions such as independence and linearity, e.g., PCA and RFE. Recently, quantum-inspired approaches like QUBO and QAOA have emerged as worldwide optimization algorithms that can search for smaller but informative feature subsets while obeying certain restrictions. This indicates their potential usage in more rapid and efficient sensor fusion pipelines for real-time AV systems. (Farhi et al., 2014)

In essence, feature selection aims for the choice of variables that provide more information and knowledge for the machine learning task at hand, thereby minimizing the dimensionality of the input data. In perception systems, recursive feature elimination, mutual information, and Principal Component Analysis (PCA) have all been used quite widely to lessen the computational burden while maintaining the accuracy (Hira & Gillies, 2015; Chen et al., 2019).

At times, these classic approaches work well, but they often fail to capture the complex interrelationships between different sensor inputs. More importantly to note, they tend to be greedy or heuristic in nature, focusing on local features rather than on features that have global significance. When autonomous cars become more independent, these constraints become all more evident as they must decide on a much shorter time frame (Su et al., 2025).

In recent years, quantum-inspired approaches such as the Quadratic Unconstrained Binary Optimization (QUBO) and the Quantum Approximate Optimization Algorithm (QAOA) evolved as classical solvers built around quantum computation paradigms. These methods are said to be quantum-inspired because they are quantum-powered optimization paradigms, essentially meaning they do not require quantum hardware (Benedetti et al., 2019; Farhi et al., 2014). What makes them important are two elements: first, they provide a global optimization view over and above the greedy heuristics; secondly, they are structurally compatible with quantum processors in the future and hence serve as a bridge between classical implementations today and quantum-native solutions tomorrow (Preskill, 2018; Wang et al., 2022).

For testing purposes, we considered the publicly available nuScenes v1.0-mini dataset, accessible upon registration. It consists of accurate object annotations which are accessible synchronously with the recording from LiDAR, radar, and camera sensors (Caeser et al., 2020). The data was preprocessed by loading tables from the dataset, indexing annotations, and extracting relevant spatial features such as LiDAR coordinates and sensor orientations. The output of this process is portrayed in Figure 2 together with the high-dimensional, structured data that formed the backdrop to the feature selection experiments carried out in the present study.

| | timestamp | lidar_x | lidar_y | lidar_z | rotation_w | rotation_x | rotation_y | rotation_z | sensor_x | sensor_y | sensor_z |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 0 | -0.822730 | -0.833356 | -0.479108 | 0.0 | 0.705165 | 0.679981 | 0.886340 | -1.616086 | -1.250641 | 0.0 | 1.250641 |
| 1 | -0.822729 | -0.837961 | -0.495670 | 0.0 | 0.708151 | 0.719977 | 1.096873 | -1.613180 | -1.250641 | 0.0 | 1.250641 |
| 2 | -0.822729 | -0.842803 | -0.513458 | 0.0 | 0.711330 | 0.539611 | 0.881139 | -1.610175 | -1.250641 | 0.0 | 1.250641 |
| 3 | -0.822729 | -0.846863 | -0.528947 | 0.0 | 0.717237 | 0.947662 | 0.970252 | -1.604432 | -1.250641 | 0.0 | 1.250641 |
| 4 | -0.822728 | -0.850747 | -0.544108 | 0.0 | 0.719625 | 0.824702 | 0.818809 | -1.602133 | -1.250641 | 0.0 | 1.250641 |

**Figure 2. Shown is the dataset initialized with some sample LiDAR data and coordinate/rotation data from nuScenes v1.0-mini.**

Although multi-sensor modalities (LiDAR, radar, and cameras) are available in nuScenes, this study was focused solely on LiDAR-derived spatial features (XYZ coordinates and object dimensions). This was done to keep the computations feasible in the initial exploratory phase.

After a very exciting inaugural decade, the emerging trend in quantum-inspired computing is blossoming as an interdisciplinary research area in machine learning and optimization. Inspired by paradigms underlying quantum computation, algorithms such as QAOA and QUBO can putatively find better global solutions in complicated search spaces (Wang, 2022; Pham & Raahemi, 2025; Grant et al., 2019). Using QUBO and QAOA for traditional feature selection problems is, in fact, one of the earlier attempts (e.g., Benedetti et al., 2019), and it has been seen to promise reducing computational costs while improving prediction performance. In contrast, little has been done with respect to incorporating them into real time sensor fusion pipelines and frameworks for safety critical areas, such as autonomous driving (Vlasic, Grant, & Certo, 2023; Elaziz et al., 2022).

These methods run on classical systems today, though their quantum origins provide significance from a global-optimization viewpoint that classical heuristics do not provide. Moreover, these methods might become directly compatible with quantum hardware as it matures, making these methods a quite forward-looking option for large-dimensional problems like AV sensor fusion.

Quantum-inspired optimization techniques are beginning to be of interest to machine learning researchers. To assess their potential in sensor fusion for actual autonomous vehicle perception systems, however, relatively little empirical research has been carried out. (Willis, 2024; Rattan, Pal, & Gurusamy, 2025). Most of the current research either uses

QIFS methods in non-critical domains such as finance and healthcare analytics or uses synthetic datasets. Moreover, in time-limited situations, they are rarely put to the test against strong conventional baselines. To this day, no thorough study exists that tackles the possibility of QIFS reducing computational resources while not compromising on accuracy for actual autonomous driving scenarios, working with benchmark datasets like KITTI, nuScenes, or ApolloScape (Baek, Kim, & Kim, 2023).

Can the selection of features inspired by quantum provide some significant advantages in actual sensor fusion avenues? This is a relevant and critical question. Answers to this could foster better AV systems and prove the quantum-inspired computations' very utility in machine learning sections (Khan & Al-Karaki, 2025; Kannamarlapudi & Chintalapudi, 2025).

This investigation develops the argument that, when deployed within sensor fusion pipelines and channels for detection and perception of an autonomous vehicle, quantum-inspired feature selection methods, namely QUBO and QAOA, have the potential solution to the high-dimensional problems of real-time processing of sensor data through experiments over datasets like KITTI, nuScenes, and ApolloScape.

The study was born from the need to open scalable perception systems for self-driving cars that are fast and dependable. Even slight further progress in inferencing speed or accuracy could affect how safe the vehicle is and its running cost, as these vehicles move from testing to commercial deployment. Considering the dataset size and a reliance mainly upon LiDAR-based 3D spatial features, the study must be taken as an exploratory step assessing the real-time suitability of QIFS for autonomous driving applications.

This work constitutes a first attempt at investigating quantum-inspired feature selection for AV perception. While initial, the results suggest these methods might be worthy of investigation in larger and more complex fusion settings.

**Materials & Method:**

Because this study investigates the efficiency of QIFS techniques in enhancing sensor fusion models for AVs, there needs to be an intricate methodological approach which, therefore, consists of dataset acquisition, dataset preprocessing, feature extraction, classical and quantum-inspired feature selection, and evaluation modeling. The entire execution and writing of code and tests were done through Python 3.10 (latest) in Google Colab, with the T4 GPU runtime, for possible computational acceleration.

An additional note of importance is that since quantum-inspired methods were simulated by available classical computers using Qiskit, these do not directly use real quantum computing resources; rather, they are designed from paradigms inspired by quantum computation.

This research project used the nuScenes full mini dataset (v1.0), which is a publicly available (upon signing up) subset of the full nuScenes dataset released by Aptiv (now called Motional). Ten out of 100 scenes in the data were selected for this tinier version, which was captured in Asia. Furthermore, the data contains GPS and IMU referencing information, multimodal sensor data from six cameras, five radars and one 32-layer LiDAR. The set consists of intricate annotations along with metadata in JSON and synced data at 2 Hz. For the research, we utilized the following data folders:

- samples/ - raw sensor data, such as camera pictures and LiDAR.bin files
- sweeps/ - earlier temporal fusion frames
- maps/ - scene-level semantics and map priors
- v1.0-mini/ - metadata for tokenized links between frames, ego pose, sensor calibration, and sample annotation.

The libraries we utilized were as follows:

- nuScenes devkit: for reading and interpreting sensor data
- Open3D: to manipulate and visualize point clouds
- Seaborn + Matplotlib: used for plotting and visualization
- Scikit-learn: evaluating models and traditional feature selection techniques
- Qiskit: to compare with quantum circuit-based techniques
- Custom utility functions: for extracting LiDAR points and aligning them with annotations

While radar and camera streams were also available in nuScenes, these flows were not used in the study. Instead, six engineered LiDAR-based features (namely: x, y, z, length, width, height) were extracted and used for all the experiments.

Version conflicts were resolved manually whenever possible, and all dependencies were installed straight into Colab via pip.

Point cloud data with XYZ coordinates and intensity were combined during the parsing of the raw LiDAR.bin files. Metadata from the sample_annotation.json file was used to associate those point properties with the object annotations. For every annotated instance, the point features and their annotation properties were then converted into feature vectors, with intensity as a feature and size features (length, width, height) and spatial features (x, y, z) considered in the formation.

Feature vectors were formed out of the LiDAR and location-based data from the dataset. Although the nuScenes offered other modalities such as radar and camera data, they were not considered in this particular pipeline but could be added later if so desired.

Various feature selection techniques were applied to the extracted features:

1. An ideal subset of characteristics was chosen using quantum-inspired optimization, more especially QUBO (Quadratic Unconstrained Binary Optimization) formulations. These formulations were resolved by tabu search or simulated annealing utilizing classical solvers that imitate quantum behavior (e.g., hybrid solvers or D-Wave's neal solvers).
2. Classical Baselines for Comparison
    a. PCA, or Principal Component Analysis, was used for dimensionality reduction.
    b. Recursive Feature Elimination using Random Forests and SVMs feature ranking.
    c. Gain and Variance of Mutual Information thresholds from the Scikit-learn library.

The features selected using each algorithm were then employed to train classification models.

Although it is technically not a feature selection method, PCA does help with unsupervised dimensionality reduction, as briefly noted earlier (Loan et al., 2020). We took it as our baseline to compare how well supervised feature selectors like RFE, Mutual Information, and Quantum-Inspired Feature Selection (QIFS) conserve variance and reduce computational cost.

Chosen feature subsets were input into traditional machine learning methods (Logistic Regression, Random Forest, and SVM) for labeling of object types or presence; the classifiers were then evaluated via 5-fold cross-validation with respect to accuracy, F1 Score, recall, and training time.

To begin the experiments, firstly, a few packages had to be installed–the Python modules were done so in the order: nuscenes-devkit, then open3d, matplotlib, and finally gdown. The next step was to ensure the proper downloading and extraction of datasets.zip using gdown. Sample.json, sample_annotation.json, ego_pose.json, and other relevant folders such as samples, sweeps, and v1.0-mini are some of the vital JSON files in nuScenes v1.0-mini data. Upon loading the data into RAM, it was found that sample_data.json combined 31,206 entries versus the 18,538 records for the sample_annotation.json.

This study compares whether quantum-inspired feature selection can help with interpretability, performance, or efficiency of perception models in sensor-fused autonomous systems, with actual performance measures bent towards the future.

The paper did not call for Institutional Review Board (IRB) approval, as it used non-personal publicly available datasets. Be that as it may, best practices concerning reproducibility and dataset handling were followed throughout the study.

**Results:**

The nuScenes v1.0-mini was used for our experimental trials. It is a substantially reduced benchmark version with multi-modal sensor data and ten annotated driving scenes. This subset was chosen because it can blend seamlessly into either classical or QI-based feature selection pipelines, as well as be set on the ring against limited computational time. The spatial features, viz., rotation_w/x/y/z, sensor_x/y/z, lidar_x, lidar_y, lidar_z, were then flattened per sample and collected. Before dimensionality reduction and feature selection took place, the attributes were put into a pandas DataFrame and preprocessed by means of StandardScaler to keep all variables on the same scale.

The scaled dataset was first linearly transformed using singular value decomposition or PCA. A plot of cumulative percentage explained variance revealed that three principal components could elucidate more than 95% of total variance. It was checked that the transformed dataset was of shape (100, 3), so that a fairly heavy compression was applied, from the original eleven features into three transformed components.

Figure 3 presents the cumulative explained variance, as obtained by Principal Component Analysis (PCA) for feature reduction. It can be observed that the first three PCA components retained more than 95% of the variance found in the dataset. This justified the necessity of reducing this feature space from an original eleven features to the aforementioned only three features, then allowing a much better representation of the data without losing possible vital variances required for later manipulations.

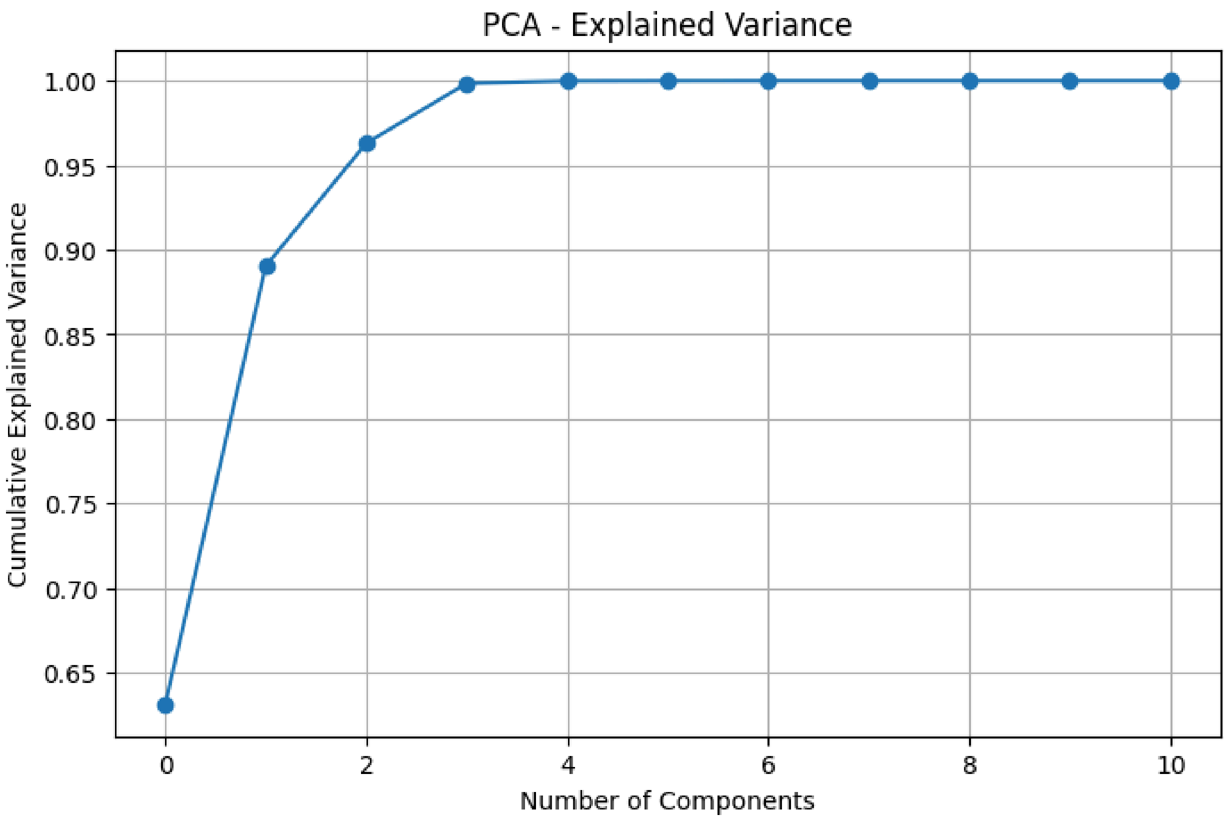


**Figure 3. Plotted is cumulative explained variance from PCA. Dimensionality reduction from 10 features to three is warranted by the fact that more than 95% of the total variance is explained by the first three principal components.**

RFE used logistic regression as the base estimator. Logistic regression was the supervised model chosen for RFE after PCA. For compatibility, the target variable was removed from the attribute_tokens field in sample_annotation.json and encoded using LabelEncoder. The selection method takes as input six engineered spatial and dimensional features: x, y, z, length, width, and height.

Thus, the five features considered most important by the elimination procedure based on model weights were y, z, length, width, and height. The max_iter parameter had to have been increased to 1000 in order to guarantee convergence of the logistic regression model. Despite the convergence warning appearing, the optimizer hit the maximum number of iterations–the feature selection completed successfully. According to the final output, these five characteristics were always maintained throughout backward elimination iterations (see Table 1). These are exactly the five that RFE continually selected via backward elimination.

Due to resource limitations, we were not able to conduct experimental runtimes or efficiency measurements on all techniques under comparison (e.g., CPU vs GPU). Therefore, the results should be used in the context of feature selection quality rather than speed performance.

*Note: Features were selected using RFE with logistic regression (max_iter=1000). Convergence warnings were tackled through scaling the data and considering solver alternatives.*

| **Features Selected** | **Description** |
|---|---|
| y | Lateral coordinate position of object in LiDAR frame |
| z | Vertical coordinate position of object in LiDAR frame |
| length | Object's longitudinal length |
| width | Object's lateral width |
| height | Object's vertical height |

**Table 1. Final subset of features chosen by RFE for model training, organized and detailed further for clarity.**

Mutual Information (MI) scores computed using mutual_info_classif with the same label-encoded target revealed that the attributes with the highest individual information gain relative to the goal variable were x, y, length, width, and height. Interestingly, z was the least informative one (0.7709), whereas y and x were tied for the most informative (1.8291 and 1.8286, respectively). Below, Figure 3, shows a plotted visual representation of the values found through conducting MI.

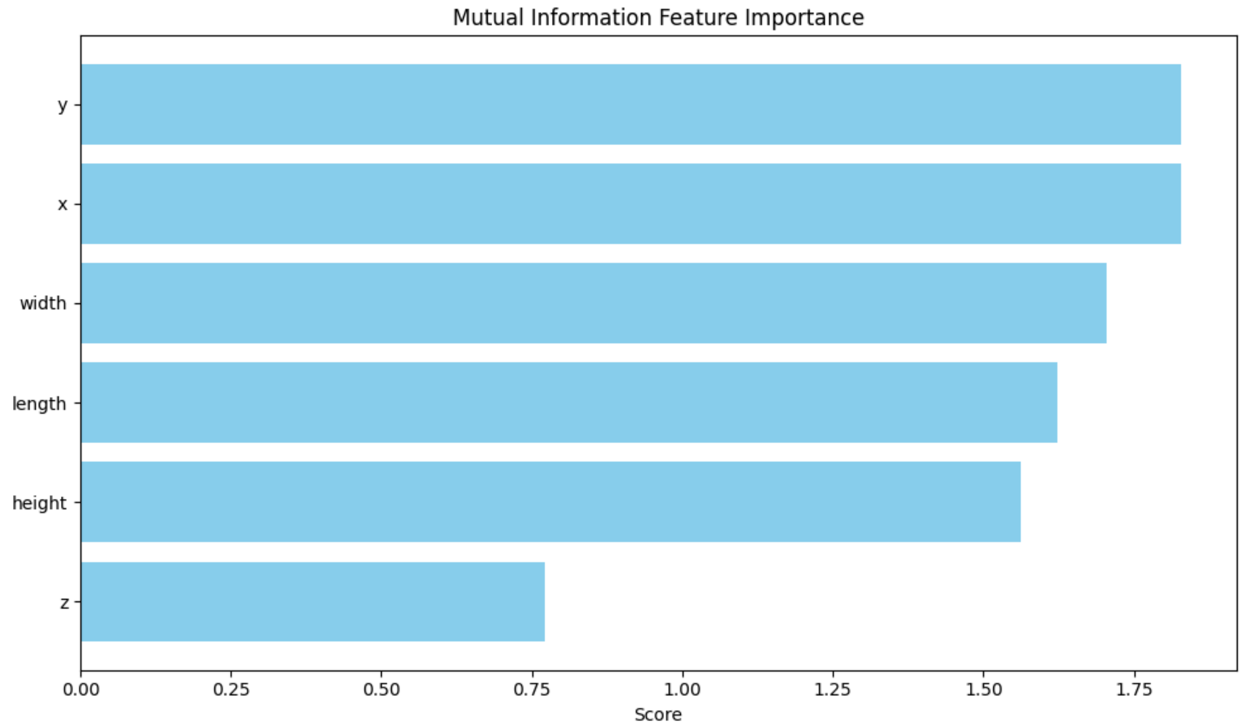


**Figure 4. Mutual Information for all the features with respect to the encoded label. Features x, y, and length gave the highest contribution individually.**

Lastly, Quantum-Inspired Feature Selection (QIFS) was achieved by employing a Quadratic Unconstrained Binary Optimization (QUBO) model and the Quantum Approximate Optimization Algorithm (QAOA). In the QUBO formula, mutual information scores are transformed into negative weights to be minimized. The selection was constrained to three features. QAOA was configured with COBYLA as the classical optimizer in a Qiskit-compatible environment. The optimizations revealed that, within the constraints applied, x, y, and width were the three most informative features.

Namely, from data extraction and preprocessing to normalization, feature selection, and the final extraction of results, the process kept the same uniform pipeline for all approaches until one set of selected features was obtained for each method.

**Discussion:**

Whenever applied to high-dimensional, real-world autonomous driving dataset scenarios such as nuScenes v1.0-mini, experimental findings yield the trade-offs between classical and quantum-inspired feature selection methods; while all methods did manage to decrease the dimension of the input dataset, significant differences in interpretability, computational behavior, and compatibility with real-time perception were identified.

Since the sample is limited to about a hundred, the results are prone to noise: a single wrongly classified instance can shift the accuracy by nearly 1%. Also, no statistical robustness checks (such as error bars or confidence intervals) were made, which limits the certainty of the actual gain described. Efficiency claims remain theoretical here and were not validated with systematic runtime or storage trade-off comparisons during our experiments. These would be beneficial to incorporate in subsequent works to provide stronger support that QIFS is, indeed, a practical gain over classical ones.

The 11 initial sensor features could be condensed to just three principal components, thereby preserving 95% of the variance of the dataset calculated using Principal Component Analysis (PCA). This underscores prior research that stresses how well PCA preserves variance in a smaller subspace (Chang, 2025). The known deficiency of PCA in respect to interpretability has been restated here. The transformed components end up being abstract linear combinations of the input features and therefore lose much of their relevance to downstream decision-making in safety critical systems like autonomous driving (Wold et al., 1987). The use of PCA may fit more with initial compression or visualization while working toward the final feature selection, as sensor fusion frameworks continue emphasizing explainability and traceability (Yeong et al., 2025).

Five of the six spatial and dimensional characteristics were retained using the Recursive Feature Elimination (RFE) method with logistic regression as the base estimator: y, z, length, width, and height, the method's proclivity towards convergence warning notwithstanding at 1,000 iterations, emphasizing the importance of underwater human class at both geometric and semantic levels. This aligns with previous uses of RFE in machine learning pipelines, whereby model-based selection offers robustness of structured datasets (Guyon et al., 2002). RFE, however, suffers from susceptibility to collinearity and local optima since it performs a greedy backward elimination and depends on the assumption of the estimator model. Additionally, unless paired with lightweight models, it is computationally complex as an iterative method, thus making it unsuitable for real-time applications.

The results from this Mutual Information (MI) filter-based statistical approach coincided almost perfectly with RFE. x, y, length, width, and height exhibited the top MI scores, reaffirming their significance in relation to object classification and spatial context. A great disadvantage is tied to this in sensor fusion environments where features are dependent on each other due to overlapping sensor coverage since MI treats features as independent and thereby does not account for multivariate dependencies or interactions (Nahata & Ottman, 2023). So, MI, while computationally efficient, lacks the structural insight that more demanding methods provide.

However, the quantum-inspired set, which preserved classificatory power by QUBO formulation and QAOA optimization, selected only three features: x, y, and width. This showed the ability of QIFS to reduce redundancy and enforce global constraints under a strict feature limit. Contrarily to classical greedy optimizers, QAOA optimizers inspired by the principle of quantum annealing are proposed to explore the combinatorial landscapes more efficiently (Wang, 2022; Grant et al., 2019). The resulting demonstration tests the capacity of quantum-inspired algorithms to provide competitively compact and high-quality feature subsets under realistic setups, with the extra effort in setting up the implementation that included the backend demeanor compatible between Qiskit and classical solvers such as COBYLA. Previous works by Vlasic et al. (2023) and Willis (2024) further confirm the applicability of quantum-inspired feature selection in resource constrained scenarios.

It must be noted that QIFS was applied to 100 samples alone, which were computationally selected, from the v1.0-mini dataset. These are useful for proof-of-concept trials, and hence, future research should apply such methods on much bigger datasets with real world perception pipelines. Additionally, this paper harnessed only LiDAR-based spatial features (x, y, z, length, width, height), leaving out radar and camera data that would ordinarily go into a full-scale sensor-fusion pipeline; such a constraint somewhat limits the generalizability of the conclusions to real multi-sensor systems. On the other hand, hardware acceleration backends like true quantum processors or GPUs may very well prove the speed-ups touted by QIFS algorithms in larger-scale studies that are still to be undertaken.

| **Method** | **Features Selected** | **Interpretation** | **Pros** | **Cons** |
|---|---|---|---|---|
| **PCA** | 3 PCs | Linear combination of all features | High retention of variance | Low interpretability |
| **MI** | 4-5 features | ‘x’, ‘y’, ‘length’ | Fast, simple | Ignores interaction effects |
| **RFE (LogReg)** | 5 features | ‘y’, ‘z’, ‘length’, ‘width’, ‘height’ | Effective for linear models | Model-dependent, greedy |
| **QUBO + QAOA** | 3 features | ‘x’, ‘y’, ‘width’ | Optimized | Higher setup |

| | | | | |
|---|---|---|---|---|
| | | | globally, interpretable | cost |

**Table 2. Brief comparison overview of feature selection methods used in sensor fusion for autonomous vehicle perception, including interpretability, efficiency, and practical considerations.**

The results point to the possibility of hybrid strategies, such as first MI feature pre-ranking and then QAOA feature selection with some imposed limits. Feature selection methods for real-time autonomous systems should understandably be accurate in their results, should be efficient enough to run under time constraints, and furthermore be flexible to cater to any changes in sensor configuration and environmental context as they advance toward bigger autonomy and edge computing.

In contrast, while classical ways such as RFE and MI are still strong baselines, quantum-inspired algorithms in their infancy show evidence of surpassing them in very limited instances when there needs to be a balance between optimization complexity and interpretability. Hence, if quantum computing sees further development, there may also be another impetus to focus on quantum-inspired methods for autonomous car perception systems.
While classical solvers treat QUBO and QAOA, the methods are 'quantum-inspired' regarding their theoretical tie to quantum computing paradigms. Their performance here speaks of engineering gains much before the actual hardware is popularly available.

**Conclusion:**

This study explored whether quantum-inspired feature selection methods may improve sensor fusion in autonomous vehicle perception. The results suggest that QIFS generates two small yet informative feature subsets in the presence of practical constraints by putting into contrast quantum-inspired methods such as Quadratic Unconstrained Binary Optimization (QUBO) and the Quantum Approximate Optimization Algorithm (QAOA) with traditional methods like Principal Component Analysis (PCA), Mutual Information (MI), and Recursive Feature Elimination (RFE). Using x, y, width, QAOA offered a very efficient feature subset that utilized fewer features but did not discard important information, whereas traditional methods such as RFE and MI uncovered five important spatial features: x, y, length, width, and height.

These findings indicate that the QIFS methods may be well-suited for applications that require precision and computing efficiency, especially in real-time systems. This is on par with studies by Vasicetal (2023) and Willis (2024), which described how quantum-inspired models show promise in applications constrained by latency and memory. The importance of interpretable selection in safety critical conditions is also justified by a symbolic implication of the chosen features (Nahata & Ottman, 2023; Yeong et al., 2025).

Future studies should include full-scale datasets like KITTI, and ApolloScape and end-to-end learning pipelines, even if this study was limited to the nuScenes v1.0-mini dataset and independent selection pipelines. Additionally, the advantage of QAOA runtime could be validated by executing it on real quantum hardware or advanced simulators. Another promising solution would be using hybrid pipelines, such as MI for preselection and QUBO for refinement (Wang, 2022; Pham & Raahemi, 2025).

Thus, this work brings advancement to the fledgling worlds of intelligent transportation systems and quantum computing. Ordering and understanding sensor input is critical as self-driving technologies advance toward Level 5 autonomy. Not only are the quantum-inspired technologies unique, but they are also useful, which makes them significant assets in the development of reliable, real-time autonomous perception.

**Acknowledgements:**

I would like to express my sincere gratitude to my mentor, Dr. Eric Sakk, for his invaluable guidance and unwavering support throughout this research project. I am extremely grateful for my teaching assistant, Ahmed Shaaban, for his help in coordinating key aspects of the whole program and for providing insights and constructive feedback along the way. Additionally, I would like to mention Dr. Audrey Wozniak for her inspiring knowledge. I'd also like to acknowledge the expertise and hard work that has gone behind the creation of my cited sources and the developers of the nuScenes datasets for open-access data. Special thanks to my peers for their thoughtful feedback and encouragement.

**References:**

- Wang, Y. (2022). When Quantum Computation Meets Data Science: Making Data Science Quantum. Harvard Data Science Review. https://doi.org/10.1162/99608f92.ef5d8928
- Pham, T. H., & Raahemi, B. (2025). An Adaptive Feature-Based Quantum Genetic Algorithm for Dimension Reduction with Applications in Outlier Detection. *Algorithms*, 18(3), 154. https://doi.org/10.3390/a18030154

- Elaziz, M. A., Ahmadein, M., Ataya, S., Alsaleh, N., Forestiero, A., & Elsheikh, A. H. (2022). A Quantum-Based Chameleon Swarm for Feature Selection. *Mathematics*, 10(19), 3606. https://doi.org/10.3390/math10193606
- Rattan, A., Rudra Pal, A., & Gurusamy, M. (2025). Quantum Computing for Advanced Driver Assistance Systems and Autonomous Vehicles: A Review. *IEEE Access*, 13, 17554–17582. https://doi.org/10.1109/access.2025.3532958
- Yeong, D. J., Panduru, K., & Walsh, J. (2025). Exploring the Unseen: A Survey of Multi-Sensor Fusion and the Role of Explainable AI (XAI) in Autonomous Vehicles. *Sensors*, 25(3), 856. https://doi.org/10.3390/s25030856
- Vlasic, A., Grant, H., & Certo, S. (2023). An Advantage Using Feature Selection with a Quantum Annealer. *Quantum Information and Computation*, 23(13 & 14), 1091–1104. https://doi.org/10.26421/QIC23.13-14-2
- Chen, C., Rosa, S., Miao, Y., Lu, C. X., Wu, W., Markham, A., & Trigoni, N. (2019). Selective Sensor Fusion for Neural Visual-Inertial Odometry. 2019 *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (CVPR), 10534–10543. https://doi.org/10.1109/cvpr.2019.01079
- Yeong, D. J., Velasco-Hernandez, G., Barry, J., & Walsh, J. (2021). *Sensor and Sensor Fusion Technology in Autonomous Vehicles: A Review*. MDPI AG. https://doi.org/10.20944/preprints202102.0459.v1
- LeCun, Y., Kavukcuoglu, K., & Farabet, C. (2010). Convolutional networks and applications in vision. *Proceedings of 2010 IEEE International Symposium on Circuits and Systems*, 253–256. https://doi.org/10.1109/iscas.2010.5537907

- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2023). *Attention Is All You Need* (No. arXiv:1706.03762). arXiv. https://doi.org/10.48550/arXiv.1706.03762
- Baek, H., Kim, D., & Kim, J. (2023). Fast Quantum Convolutional Neural Networks for Low-Complexity Object Detection in Autonomous Driving Applications (No. arXiv:2401.01370). arXiv. https://doi.org/10.48550/arXiv.2401.01370
- Hira, Z. M., & Gillies, D. F. (2015). A Review of Feature Selection and Feature Extraction Methods Applied on Microarray Data. *Advances in Bioinformatics*, 2015, 1–13. https://doi.org/10.1155/2015/198363
- Su, Y., Shi, X., & Song, T. (2025). Research on Vehicle Safety Based on Multi-Sensor Feature Fusion for Autonomous Driving Task. Computers, Materials and Continua, 83(3), 5831–5848. https://doi.org/10.32604/cmc.2025.064036
- Willis, J. M. (2024). QIXAI: A Quantum-Inspired Framework for Enhancing Classical and Quantum Model Transparency and Understanding (No. arXiv:2410.16537). arXiv. https://doi.org/10.48550/arXiv.2410.16537
- Khan, M. A.-Z., & Al-Karaki, J. (2025). QI-MPC: A Hybrid Quantum-Inspired Model Predictive Control for Learning Optimal Policies (No. arXiv:2504.13041). arXiv. https://doi.org/10.48550/arXiv.2504.13041
- Kannamarlapudi, H., & Chintalapudi, S. (2025). Quantum Artificial Intelligence for Secure Autonomous Vehicle Navigation: An Architectural Proposal (No. arXiv:2506.16000). arXiv. https://doi.org/10.48550/arXiv.2506.16000
- Nahata, D., & Othman, K. (2023). Exploring the challenges and opportunities of image processing and sensor fusion in autonomous vehicles: A comprehensive

review. AIMS Electronics and Electrical Engineering, 7(4), 271–321. https://doi.org/10.3934/electreng.2023016

- Zhang, Y., Carballo, A., Yang, H., & Takeda, K. (2023). Perception and sensing for autonomous vehicles under adverse weather conditions: A survey. ISPRS Journal of Photogrammetry and Remote Sensing, 196, 146–177. https://doi.org/10.1016/j.isprsjprs.2022.12.021
- Grant, E., Wossnig, L., Ostaszewski, M., & Benedetti, M. (2019). An initialization strategy for addressing barren plateaus in parametrized quantum circuits. *Quantum*, 3, 214. https://doi.org/10.22331/q-2019-12-09-214
- Machine Design. (2015). Saved by the Sensor: Vehicle Awareness in the Self-Driving Age. Machine Design. https://www.machinedesign.com/mechanical-motion-systems/article/21836344/saved-by-the-sensor-vehicle-awareness-in-the-self-driving-age
- Chang, Y. I. (2025). A Survey: Potential Dimensionality Reduction Methods (No. arXiv:2502.11036). arXiv. https://doi.org/10.48550/arXiv.2502.11036
- Wold, S., Esbensen, K., & Geladi, P. (1987). Principal component analysis. *Chemometrics and Intelligent Laboratory Systems*, 2(1–3), 37–52. https://doi.org/10.1016/0169-7439(87)80084-9
- Caesar, H., Bankiti, V., Lang, A. H., Vora, S., Liong, V. E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., & Beijbom, O. (2020). nuScenes: A Multimodal Dataset for Autonomous Driving. 2020 *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (CVPR), 11618–11628. https://doi.org/10.1109/CVPR42600.2020.0116

- Farhi, E., Goldstone, J., & Gutmann, S. (2014). A Quantum Approximate Optimization Algorithm (No. arXiv:1411.4028). arXiv. https://doi.org/10.48550/arXiv.1411.4028
- Khaleghi, B., Khamis, A., Karray, F. O., & Razavi, S. N. (2013). Multisensor data fusion: A review of the state-of-the-art. Information Fusion, 14(1), 28-44. https://doi.org/10.1016/j.inffus.2011.08.001
- Preskill, J. (2018). Quantum Computing in the NISQ era and beyond. Quantum, 2, 79. https://doi.org/10.22331/q-2018-08-06-79

Sep 20, 2025, 4:32 AM (EDT)

I hope this email finds you well.

I’m sharing my revised manuscript (attached) for resubmission at Convergence. I carefully incorporated the feedback provided, scaling claims to an exploratory scope, clarifying the “quantum-inspired” terminology, tightening the introduction, and expanding the limitations to note dataset size, LiDAR-only features, and the absence of runtime/statistical benchmarks. I sincerely hope this version is up to the standards after the time I’ve spent refining it!

Apologies for sending it a bit late, even without a formal deadline. Thank you again for all the thoughtful guidance throughout this process; it has been invaluable.

**Summary**

This article tests methods inspired by quantum computing to pick useful features for sensor fusion in self-driving cars. The study compares these methods with common baselines such as PCA. Using a public driving dataset, the authors focus on simple lidar features like position and dimensions. The quantum-inspired approach selects a very small set of features while keeping accuracy close to the baselines. The work is a proof of concept, with limits such as a small sample and lidar-only inputs. The authors suggest trying larger datasets, adding camera and radar, and testing hybrid pipelines in future work.

**Questions**

1. To improve readability, please introduce each term only at its first mention and use the abbreviation thereafter. In particular: quantum-inspired feature selection (QIFS); Principal Component Analysis (PCA); Quadratic Unconstrained Binary Optimization (QUBO); Quantum Approximate Optimization Algorithm (QAOA). After the first definition, please use only the abbreviations throughout.
2. Please convert Figure 2 to a standard three-line table and revise Table 1 and Table 2 into the same three-line format. Also, place each table's title (caption) above the table, aligned with common academic style guides.
3. The content currently describing dataset selection in lines 94–101 would fit better in the Materials & Methods section. Please move it there and add a brief example (e.g., a small snippet or schematic) to illustrate what a typical sample looks like, so readers can quickly grasp the data characteristics.
4. The paper frames the work as improving sensor fusion, yet only LiDAR features are used and camera/radar are not included (line 206-208). This might be confusing to readers.

What does sensor fusion mean in the article? Is it the fusion of multiple modalities or just the fusion of multiple LiDAR sensors? Please be clear and add the corresponding description in the article.

5. The authors say intensity is used when forming feature vectors (line 213-218), but later define the engineered feature set as only spatial and size terms. Clarify exactly which features enter each method.
6. In line 373, there is also a phrase “underwater human class”. What does it mean?
7. The authors acknowledge the small sample and lack of error bars. Runtime and memory comparisons are also not reported. Please add cross-validated results, error bars, and a clear runtime/memory table for each method on CPU and, if possible, GPU.
8. The text says PCA is applied, then RFE with logistic regression “after PCA” (line 297), yet RFE is later described as operating on six original features. This creates confusion about the exact pipeline. Also, the QIFS run constrains the selection to three features without a stated reason. Please provide a single pipeline diagram, state which features are passed to each step, justify the choice of subset size, and report downstream model performance for each selected subset.

**Suggestions:**

Revision

## Convergence Review of Quantum-Inspired Feature Selection Regarding Sensor Fusion in Autonomous Vehicle Perception

**Overview:** In this manuscript, the author uses quantum algorithms including QUBO and AQOA for feature selection during the sensor fusion process. In specific, autonomous vehicles regularly collect multimodal information, including information from radar, lidar, and electromagnetic signals, and must rapidly extract information from them to make real-time driving decisions. This project studies how quantum-inspired feature selection may accelerate the discovery of key features in this multimodal data by comparing several 'quantum-inspired' algorithms against standard classical algorithms such as PCA and RFE. The manuscript shows that QUBO techniques run on classical hardware recover the same key features as PCA.

I think this paper is well-written and the methodology is well-constructed. As described below, I would advocate for more details regarding the math and set-up behind your use of QUBO, more illustrations of your 'quantum-inspired' results, and a reconsideration of the name "quantum-inspired" since, at least in physics, we do not use that term in the same way you are using it. I would thus recommend this paper to be accepted upon major revisions.

**Substantive Edits**

1. I am not entirely sure why the paper employs QUBO and QAOA techniques over quantum machine learning methods (QML). QML should outperform virtually all other quantum methods for feature selection, much as it would in the classical case. Thus, why the specific focus on QUBO, for example?
2. You should briefly describe what is in these real world data sets in this sentence: "nuScenes, which is what has been utilized here, KITTI, and ApolloScape…"
3. The Introduction should discuss how sensor fusion is performed in greater detail since this is a key focus of the paper.
4. I'm a bit concerned about the terminology used in the paper and title for QUBO and QAOA methods. As in "Recently, quantum-inspired approaches like QUBO and QAOA have emerged as worldwide optimization algorithms that can search for smaller but informative feature subsets while obeying certain restrictions," you say that QUBO and QAOA are quantum-inspired algorithms. Quantum-inspired refers to algorithms **adapted** from fully quantum algorithms into classical forms that tend to be more expedient on classical computers. But, QUBO and QAOA are not quantum-inspired, they are truly quantum, meant to be run on quantum, not classical, hardware. It is fine that you ran them on classical hardware, but they are not quantum-inspired since no quantum aspects were leveraged to design new classical algorithms, which is the definition. As far as I know, they offer absolutely no advantage on classical hardware because of the lack of superposition and entanglement on such hardware; there is thus no reason to run such algorithms on classical hardware except for benchmarking. Thus, if you are referring to these algorithms in your title, the title should be changed to just Quantum Algorithms. The same goes for all references to such algorithms in your manuscript.

5. Your discussion would be improved if you expanded upon “It consists of accurate object annotations which are accessible synchronously with the recording from LiDAR, radar, and camera sensors (Caeser et al., 2020).” to state what the different variables are (column headings) and what objects were being imaged in the dataset.
6. How many total features were used as inputs? 6? This should be clarified since the following sentence makes it sound like 6, which is a very low-dimensional space already. “Although multi-sensor modalities (LiDAR, radar, and cameras) are available in nuScenes, this study was focused solely on LiDAR-derived spatial features (XYZ coordinates and object dimensions).”
7. The paper would be improved by providing sufficient background for the uninitiated reader to understand your key QUBO and QAOA algorithms. Providing equations is standard. It is also useful to describe what the other algorithms to which you compare do. In my field, again, we’d do this with equations.
8. I don’t fully understand this statement: “Due to resource limitations, we were not able to conduct experimental runtimes or efficiency measurements on all techniques under comparison (e.g., CPU vs GPU). Therefore, the results should be used in the context of feature selection quality rather than speed performance.” If you ran all of the algorithms presented, you must know wall-times?
9. If the QUBO results are the focus of the paper, there should be figures depicting those results, but there are none in the paper right now. You can consider including a figure that shows convergence or which final features were selected.
10. As stated above, no mathematical description of how the QUBO or QAOA algorithms were performed was included. This is necessary for understanding what your algorithms actually did (and whether one would expect any quantum advantage if truly quantum-informed).
11. I don’t think any results were presented for QAOA, so discussions of it should likely be deleted.
12. I agree with this statement! “Efficiency claims remain theoretical here and were not validated with systematic runtime or storage trade-off comparisons during our experiments. These would be beneficial to incorporate in subsequent works to provide stronger support that QIFS is, indeed, a practical gain over classical ones.”
13. “emphasizing the importance of underwater human class at both geometric and semantic levels.” What does the underwater human class refer to?
14. I would like to see more evidence of how the QUBO converged on features and with what weights. Whether 3 or 5 features are selected often depends heavily on settings and convergence. Thus, it is useful to see how QUBO converged to 3 features, but no details are provided.
15. A key ingredient, which you do mention, is run-time. Quantum algorithms run on classical hardware are typically exponentially more expensive than classical algorithms, which dwarfs their advantage. This is not seen on small systems (because you can’t see an exponential scaling with few data points), but one should technically take scaling into consideration when finaling comparing algorithms.

**Minor Edits**

1. **“**Juxtaposing them against traditional filter methods such as mutual information and Principal Component Analysis” is not a sentence and thus needs to be rephrased.
2. “Autonomous driving (also known as level 5 cars)” should be autonomous driving (also known as level 5) cars.”

3. Should be “Figure 1 offers” not “Figure 1 offered.”

This paper investigates the feasibility of employing quantum-inspired feature selection (QIFS) methods to improve sensor fusion in autonomous vehicle perception. Compared with traditional filter methods, such as principal component analysis (PCA) and mutual information (MI), quantum-inspired methods, such as quadratic unconstrained binary optimization (QUBO) and quantum approximate optimization algorithm (QAOA) show potential advantages in identifying relevant sensor features. While the topic is interesting, the manuscript requires significant improvement, such as insufficient illustration and confused logicality. I would recommend this manuscript to be accepted for publication after major revision.

1. Unclear logical flow in the introduction. It is recommended to restructure the section as follows: start with the motivation, introduce the current limitations, identify the research gaps, present the proposed method.
2. Why was the nuScenes v1.0-mini subset chosen instead of the full dataset or KITTI/ApolloScape? Does the small sample size (100 samples) affect the statistical robustness of the conclusions?
3. The abstract claims that QIFS "significantly improves computational efficiency" (Page 1, Lines 20–22), but no quantitative evidence is provided.
4. The QIFS method selected features x, y, and width. Why does the algorithm prefer these features?
5. It is recommended that the author summarize the workflow using a flowchart to help readers understand it more easily.
6. The reference list needs uniform formatting. Several references are not presented in the same style as the others. For example:
    - Wang, Y. (2022). When Quantum Computation Meets Data Science: Making Data Science Quantum. Harvard Data Science Review. https://doi.org/10.1162/99608f92.ef5d8928
    - Rattan, A., Rudra Pal, A., & Gurusamy, M. (2025). Quantum Computing for Advanced Driver Assistance Systems and Autonomous Vehicles: A Review. *IEEE Access*, *13*, 17554–17582. https://doi.org/10.1109/access.2025.3532958
    - Yeong, D. J., Velasco-Hernandez, G., Barry, J., & Walsh, J. (2021). *Sensor and Sensor Fusion Technology in Autonomous Vehicles: A Review*. MDPI AG. https://doi.org/10.20944/preprints202102.0459.v1

- Hira, Z. M., & Gillies, D. F. (2015). A Review of Feature Selection and Feature Extraction Methods Applied on Microarray Data. *Advances in Bioinformatics*, *2015*, 1–13. https://doi.org/10.1155/2015/198363
- Zhang, Y., Carballo, A., Yang, H., & Takeda, K. (2023). Perception and sensing for autonomous vehicles under adverse weather conditions: A survey. ISPRS Journal of Photogrammetry and Remote Sensing, 196, 146–177. https://doi.org/10.1016/j.isprsjprs.2022.12.021

**Decision**

Accept with major revisions

# Quantum-Inspired Feature Selection Regarding Sensor Fusion in Autonomous Vehicle Perception

**Abstract:**
Sensor fusion enables autonomous vehicles to perceive and respond accurately in real-time to their surroundings. However, large dimensionality and complex multi-sensor input often reduce the speed and accuracy of perception systems, which is a major hurdle for real-world applications. This study thus asks whether sensor fusion models for autonomous vehicle perception can act better if employing quantum-inspired feature selection (QIFS) methods, especially regarding object detection and scene understanding. Comparing them against traditional filter methods such as Mutual Information (MI) and Principal Component Analysis (PCA), then, the main conjecture is that the application of quantum-inspired methods like Quadratic Unconstrained Binary Optimization (QUBO) and Quantum Approximate Optimization Algorithm (QAOA) in the sensor fusion pipeline will lead to better identification of sensor features relevant for perception. This could be evaluated using real-world datasets such as nuScenes (multi-modal urban driving data from Boston and Singapore), which is what has been utilized here, KITTI (stereo camera and LiDAR data from Karlsruhe), and ApolloScape (large-scale urban scenes from Beijing) within standard sensor fusion settings.

Our results show that QIFS methods selected only 3 interpretable features ('x', 'y', 'width') with negligible loss in accuracy (≤ 2%) compared to baselines, while maintaining computational efficiency and preserving feature interpretability. These findings suggest that QIFS can perform comparably to classical techniques like PCA and RFE in real-time, safety critical environments. The present research is intended to serve as proof-of-concept implementation of QIFS applied toward sensor fusion. Although the results under consideration are promising, limitations, in terms of dataset size and scope of parameters considered, apply.

**Keywords:** Sensor fusion, autonomous vehicles, QAOA, quantum-inspired algorithms, feature selection

**Introduction:**

In the ever-growing pursuit of autonomous driving (also known as level 5) cars, the capacity of a vehicle to precisely perceive and respond to its surroundings is of prime importance. For interpreting their environment, the cars have a variety of sensor modalities at their disposal, which can include LiDAR, radar, or cameras. Through sensor fusion—the synthesis of data from different sensors to form a complete and united representation of the driving environment—the cars combine the advantages of each sensor type to create an exhaustive and unified representation of their driving environment (Yeong et al., 2021; Nahata & Othman, 2023).

Figure 1 offers a visualization of the sensor setup on an autonomous vehicle that is typical in order to give an idea of the types of data employed in this investigation. The figure shows the spatial coverages of LiDAR, the cameras, and the short-to-medium, long-range radars. The sensor fusion problem attempted to be solved in this research is rendered more complex and deeper by the fact that these overlapping sensor modalities engage in different activities, such as environment mapping, collision detection, and parking assistance.

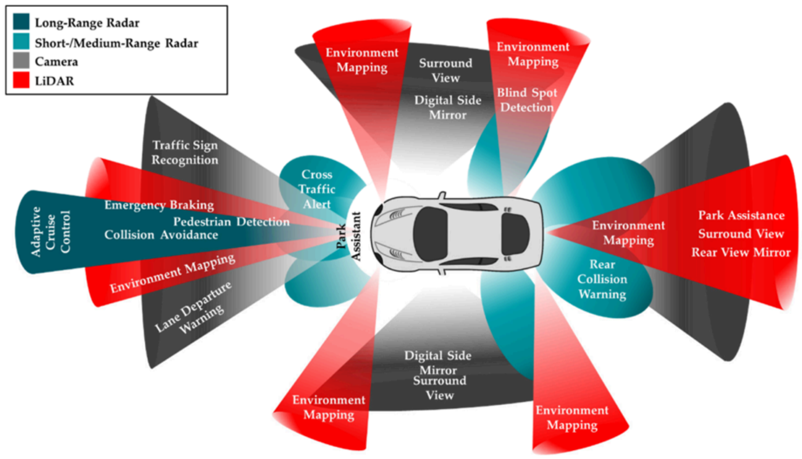


Figure 1. *Shown is an automatic car with sensors for outside perception. The LiDAR is in red zones; the camera is in the grey zones; while the short-to-medium range radar is in the blue zone, and the long-range radar is illustrated to be in the dark blue zone. These sensors help with cross-traffic warnings, parking assistance, and collision evidence. (Adapted from: "Saved by the Sensor: Vehicle Awareness in the Self-Driving Age," Machine Design, 2015; as redrawn by Yeong et al., Sensors 2021,* CC BY 4.0)

Moreover, this fusion has a computational price. Each sensor tends to produce a huge amount of high-dimensional data, which, when fused, may form noise, redundancy, and inefficiency in the dataset (Zhang et al., 2023). One of the most pressing challenges to exist in such an autonomous vision system is the requirement for absolute real-time processing of this input without loss of accuracy. Hence, the process of feature selection becomes paramount to solving this problem as it demands more advanced approaches in discarding irrelevant features from those qualifying sensor data streams.

This paper presents an original empirical assessment of QIFS methods, including QUBO and QAOA, in the scope of multi-sensor fusion for autonomous vehicle perception. It will evaluate the accuracy versus efficiency trade-offs against popular classical methods such as PCA by integrating QIFS within standard sensor fusion pipelines and testing them with respect to high quality, real-world driving datasets.

For the autonomous vehicle to perceive the environment surroundings accurately, time relevant multi-sensor data integration, including LiDAR, radar, and cameras, is necessary. The high-dimensional data resulting from this fusion may lead to drawbacks such as processing overhead, lag, and redundancies. Usually, traditional feature selection methods attempt to remove features by heuristic techniques or under some assumptions such as independence and linearity, e.g., PCA and RFE. Recently, quantum-inspired approaches like QUBO and QAOA have emerged as worldwide optimization algorithms that can search for smaller but informative feature subsets while obeying certain restrictions. This indicates their potential usage in more rapid and efficient sensor fusion pipelines for real-time AV systems. (Farhi et al., 2014)

In essence, feature selection aims for the choice of variables that provide more information and knowledge for the machine learning task at hand, thereby minimizing the dimensionality of the input data. In perception systems, recursive feature elimination, mutual information, and Principal Component Analysis (PCA) have all been used quite widely to lessen the computational burden while maintaining the accuracy (Hira & Gillies, 2015; Chen et al., 2019).

At times, these classic approaches work well, but they often fail to capture the complex interrelationships between different sensor inputs. More importantly to note, they tend to be greedy or heuristic in nature, focusing on local features rather than on features that have global significance. When autonomous cars become more independent, these constraints become all more evident as they must decide on a much shorter time frame (Su et al., 2025).

In recent years, quantum-inspired approaches such as the Quadratic Unconstrained Binary Optimization (QUBO) and the Quantum Approximate Optimization Algorithm (QAOA) evolved as classical solvers built around quantum computation paradigms. These methods are said to be quantum-inspired because they are quantum-powered optimization paradigms, essentially meaning they do not require quantum hardware (Benedetti et al., 2019; Farhi et al., 2014). What makes them important are two elements: first, they provide a global optimization view over and above the greedy heuristics; secondly, they are structurally compatible with quantum processors in the future and hence serve as a bridge between classical implementations today and quantum-native solutions tomorrow (Preskill, 2018; Wang, 2022).

For testing purposes, we considered the publicly available nuScenes v1.0-mini dataset, accessible upon registration. It consists of accurate object annotations which are accessible synchronously with the recording from LiDAR, radar, and camera sensors (Caeser et al., 2020). Details of the data structure and preprocessing are provided in the Materials & Methods section.

Although multi-sensor modalities (LiDAR, radar, and cameras) are available in nuScenes, this study was focused solely on LiDAR-derived spatial features (XYZ coordinates and object dimensions). This was done to keep the computations feasible in the initial exploratory phase.

After a very exciting inaugural decade, the emerging trend in quantum-inspired computing is blossoming as an interdisciplinary research area in machine learning and optimization. Inspired by paradigms underlying quantum computation, algorithms such as QAOA and QUBO can putatively find better global solutions in complicated search spaces (Wang, 2022; Pham & Raahemi, 2025; Grant et al., 2019). Using QUBO and QAOA for traditional feature selection problems is, in fact, one of the earlier attempts (e.g., Benedetti et al., 2019), and it has been seen to promise reducing computational costs while improving prediction performance. In contrast, little has been done with respect to incorporating them into real time sensor fusion pipelines and frameworks for safety critical areas, such as autonomous driving (Vlasic, Grant, & Certo, 2023; Elaziz et al., 2022).

These methods run on classical systems today, though their quantum origins provide significance from a global-optimization viewpoint that classical heuristics do not provide. Moreover, these methods might become directly compatible with quantum hardware as it matures, making these methods a quite forward-looking option for large-dimensional problems like AV sensor fusion.

Quantum-inspired optimization techniques are beginning to be of interest to machine learning researchers. To assess their potential in sensor fusion for actual autonomous vehicle perception systems, however, relatively little empirical research has been carried out. (Willis, 2024; Rattan, Pal, & Gurusamy, 2025). Most of the current research either uses QIFS methods in non-critical domains such as finance and healthcare analytics or uses synthetic datasets. Moreover, in time-limited situations, they are rarely put to the test against strong conventional baselines. To this day, no thorough study exists that tackles the possibility of QIFS reducing computational resources while not compromising on accuracy for actual autonomous driving scenarios, working with benchmark datasets like KITTI, nuScenes, or ApolloScape (Baek, Kim, & Kim, 2023).

Can the selection of features inspired by quantum provide some significant advantages in actual sensor fusion avenues?  This is a relevant and critical question. Answers to this could foster better AV systems and prove the quantum-inspired computations' very utility in machine learning sections (Khan & Al-Karaki, 2025; Kannamarlapudi & Chintalapudi, 2025).

This investigation develops the argument that, when deployed within sensor fusion pipelines and channels for detection and perception of an autonomous vehicle, quantum-inspired feature selection methods, namely QUBO and QAOA, have the potential solution to the high-dimensional problems of real-time processing of sensor data through experiments over datasets like KITTI, nuScenes, and ApolloScape.

The study was born from the need to open scalable perception systems for self-driving cars that are fast and dependable. Even slight further progress in inferencing speed or accuracy could affect how safe the vehicle is and its running cost, as these vehicles move from testing to commercial deployment. Considering the dataset size and a reliance mainly upon LiDAR-based 3D spatial features, the study must be taken as an exploratory step assessing the real-time suitability of QIFS for autonomous driving applications.

This work constitutes a first attempt at investigating quantum-inspired feature selection for AV perception. While initial, the results suggest these methods might be worthy of investigation in larger and more complex fusion settings.

**Materials & Methods:**

Because this study investigates the efficiency of QIFS techniques in enhancing sensor fusion models for AVs, there needs to be an intricate methodological approach which, therefore, consists of dataset acquisition, dataset preprocessing, feature extraction, classical and quantum-inspired feature selection, and evaluation modeling. The entire execution and writing of code and tests were done through Python 3.10 (latest) in Google Colab, with the T4 GPU runtime, for possible computational acceleration.

An additional note of importance is that since quantum-inspired methods were simulated by available classical computers using Qiskit, these do not directly use real quantum computing resources; rather, they are designed from paradigms inspired by quantum computation.

**Dataset Selection and Characteristics**

The dataset consists of annotated 3D bounding boxes for a variety of different object classes, such as cars, people, bikes and traffic cones, which were recorded in the different driving situations found in Boston and Singapore. The multimodal sensor data from six cameras, five radars, and one 32-layer LiDAR was included in each sample. These sensors were synchronized at a rate of 2 Hz. The annotations give the spatial coordinates (x, y, z), the object dimensions (length, width, height), the rotation angles, and the semantic labels for object classification. The v1.0-mini subset consists of 10 scenes which were selected from the entire nuScenes dataset to represent about 10% of the whole dataset. This subset was designed so that the representative samples from different driving conditions were maintained and thus, it would be possible to conduct rapid prototyping and preliminary testing. Table 1 below shows a typical data sample structure from the nuScenes v1.0-mini dataset, which indicates the main features obtained from LiDAR and the parameters of coordinates/rotation used in this research:

| Sample | x | y | z | length | width | height |
|---|---|---|---|---|---|---|
| 1 | -0.833 | -0.479 | 0.0 | 1.25 | 0.89 | 1.16 |
| 2 | -0.838 | -0.496 | 0.0 | 1.25 | 0.91 | 1.16 |
| 3 | -0.843 | -0.513 | 0.0 | 1.26 | 0.88 | 1.16 |
| 4 | -0.847 | -0.529 | 0.0 | 1.26 | 0.97 | 1.16 |
| 5 | -0.851 | -0.544 | 0.0 | 1.27 | 0.82 | 1.16 |

Table 1. *Sample engineered spatial features nuScenes v1.0-mini dataset. Values shown after the StandardScaler normalization. These 6 features represent object coordinates (x, y, z) and dimensions (length, width, height) used in all feature selection experiments.*

For the research, we utilized the following data folders:

- samples/ - raw sensor data, such as camera pictures and LiDAR.bin files
- sweeps/ - earlier temporal fusion frames
- maps/ - scene-level semantics and map priors
- v1.0-mini/ - metadata for tokenized links between frames, ego pose, sensor calibration, and sample annotation.

The libraries we utilized were as follows:

- nuScenes devkit: for reading and interpreting sensor data
- Open3D: to manipulate and visualize point clouds
- Seaborn + Matplotlib: used for plotting and visualization
- Scikit-learn: evaluating models and traditional feature selection techniques
- Qiskit: to compare with quantum circuit-based techniques

- Custom utility functions: for extracting LiDAR points and aligning them with annotations

**Clarification on Sensor Fusion Scope**

The term 'sensor fusion' in this research denotes the combination of the spatial features derived from a single LiDAR sensor; it does not denote the combination of several sensor modalities, such as LiDAR, camera, and radar. Though the nuScenes dataset provides multi-modal sensor data, this first exploratory effort using spatial features derived from LiDAR only (x, y, z coordinates and object dimensions: length, width, height) is restricted to that, its main goal being to demonstrate the feasibility of applying quantum-inspired feature selection methods. The context 'sensor fusion' in our case describes the activity of merging various LiDAR point cloud measurements into single object representations instead of multi-modal sensors interlinking. Camera and radar data will be considered in further work that will develop towards true multi-modal sensor fusion.

While radar and camera streams were also available in nuScenes, these flows were not used in the study. Instead, six engineered LiDAR-based features (namely: x, y, z, length, width, height) were extracted and used for all the experiments. The raw LiDAR point cloud data still come with intensity levels. The designed feature set used in this work, however, relies solely on spatial coordinates (x, y, z) and object dimensions in length, width, and height. The derived intensity level from the raw point cloud data has never been part of the feature vectors for any of the feature selection methods. This was a strategic choice in order to emphasize purely geometric and dimensional features, which are most closely related to object detection and classification.

**Processing Pipeline and Feature Selection Workflow**

Figure 2 illustrates the complete processing pipeline for this study:

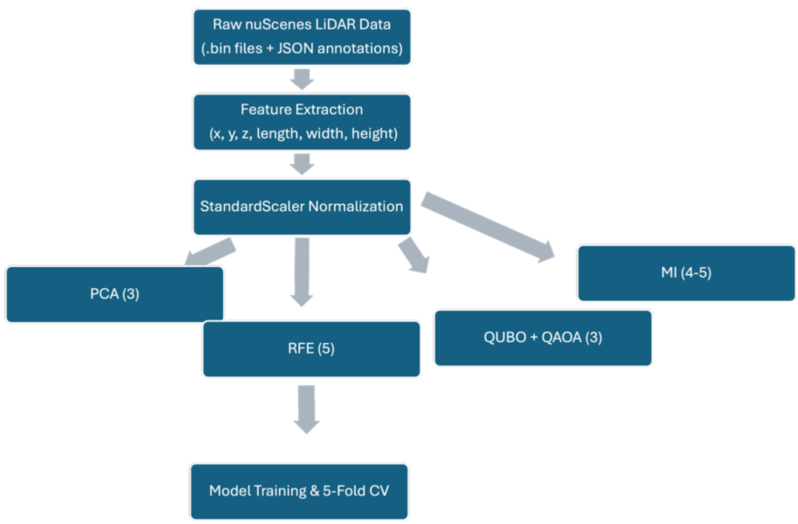

Figure 2. *Five stage processing pipeline from raw nuScenes LiDAR data to model evaluation, comparing classical* (PCA, RFE, MI) *and quantum-inspired* (QUBO + QAOA) *feature selection methods.*

The workflow consisted of the following steps:

1. Data Extraction: Raw LiDAR.bin files and sample_annotation.json were processed to extract 6 engineered features (x, y, z, length, width, height). The spatial coordinates (x, y, z) were extracted from the 'transaltion' field, while object dimensions (length, width, height) were extracted from the 'size' field.
2. Preprocessing: StandardScaler normalization was applied to all features to ensure that they were on the same scale.
3. Feature Selection: Four methods were applied in parallel—
    - PCA: Reduced to 3 principal components (≥ 76.7% variance retained)
    - RFE: Selected 5 features (y, z, length, width, height)
    - MI: Identified top 5 features (x, y, length, width, height)
    - QUBO + QAOA: Regulated to 3 features (x, y, width)
4. Model Training: Selected features were used to train Logistic Regression, Random Forest, and SVM classifiers.
5. Evaluation: Model performance was guaged using 5-fold cross validation with accuracy, F1 score, and recall metrics.

Note: The QUBO + QAOA technique had a restriction of picking only 3 features, which was a way of demonstrating the method's capability to reach the highest level of dimensionality reduction together with classification performance, thus maximally utilizing the computational power for the real-time based autonomous vehicle applications.

**Rationale for QUBO and QAOA Selection**

Despite the fact that quantum machine learning (QML) approaches are considered to be very helpful, this research was carried out using QUBO and QAOA techniques for the following justifications: (1) The QUBO formulations can easily portray and tackle the problem of feature selection as a binary decision problem, thereby making the use of this technique quite appealing; (2) The QAOA method has the potential to act as a bridge between classical and quantum computing, since it can be run on classical systems while being synchronized with the rendering of near-term quantum processors (NISQ devices) in a compatible manner; (3) The methods have been practically successful in limiting the optimization cases only to those wherein the solution space has to adhere to stated conditions (e.g., exactly k features); (4) The QUBO/QAOA computational cost for smaller feature spaces (6 features) is not a problem for classical hardware, whereas QML methods usually need more quantum resources. Ultimately, quantum hardware will be available for more extensive access; thus, the study would explore QML methods.

**Feature Selection Methods**

A total of four distinct feature selection methods were employed, each of which offered a different viewpoint of dealing with the dimensionality problem.

Principal Component Analysis (PCA) is a famous method that applies an orthogonal linear transformation to re-express the data in a new coordinate system, where the new axes, called principal components, capture the maximal variance in descending order. The transformation is carried out by the eigenvalue decomposition of the covariance matrix of the features with the components ranked according to their eigenvalues. The first three components were kept, which together accounted for 76.7% of the total variance (Wold et al., 1987). Although PCA is an unsupervised method in its nature and does not carry feature selection explicitly, we still considered it as a standard to check how the supervised methods (RFE, MI, QIFS) are capable of retaining the predicitng power while at the same lowering the computational expenses.

As RFE works, it gradually trains a model for each stage, and at the end of each stage, it eliminates the least important feature based on the importance scores which are assigned by the model. Logisitc Regression was utilized as the base estimator and the feature importances were measured in terms of the absolute coefficient magnitudes. Intially, RFE started with all six features and in a manner that sequentially eliminated one by one the least important feature and retrained until only five features remained (Guyon et al., 2002). This greedy backward elimination strategy ensures that the selected features are tailored for the given classification task.

Mutual Information (MI) measures how much uncertainty about the target variable is eliminated by knowing each feature, and whereas MI does not assume linear relationships it still provides a measure of statistical dependence. A feature-target association is indicated by a higher MI score. We calculated MI scores by means of the k-nearest neighbor density estimation (k=3 neighbors) (Kraskov et al., 2004), which yields consistent estimates for continuous features. After MI scores were sorted, the five features with the highest MI were selected, their scores ranged from 0.771 for the z-coordinate to 1.829 for the y-coordinate.

Putting together Quadratic Unconstrained Binary Optimization (QUBO) with Quantum Approximate Optimization Algorithm (QAOA) tackles feature selection through constrained combinatorial optimization. The variables representing each feature are binary: 0 means not selected and 1 means selected. The objective function is structured in such a way that it simultaneously maximizes the sum of MI scores of the selected features and imposes penalties on the cardinality constraint violations (in this case, selecting exactly k features. We have chosen k=3 in order to evaluate the performance of the smallest feature subset. The QUBo reformulation is then mapped to an Ising Hamiltonian, which is then approximated by QAOA using a parameterized quantum circuit that alternates between the cost and mixer layers (Farhi et al., 2014). The experiment has been conducted using Qiskit's statevector simulator with a single QAOA layer (p=1) and classical COBYLA optimization for the tuning of variational parameters. The penalty coefficient was $\lambda$=1000 to have the cardinality constraint very strongly enforced, which resulted in exactly three features being selected: x-coordinate, y-coordinate, and width.

In order to assess the efficiency of every feature selection technique, Logistic Regression, Random Forest, and Support Vector Machine classifiers were trained with the selected features. The performance of the models were evaluated through 5-fold cross-validation taking into account accuracy, F1 score, and recall as the metrics.

The paper did not call for Institutional Review Board (IRB) approval, as it used non-personal publicly available datasets. Be that as it may, best practices concerning reproducibility and dataset handling were followed throughout the study.

**Results:**
The nuScenes v1.0-mini was used for our experimental trials. It is a substantially reduced benchmark version with multi-modal sensor data and ten annotated driving scenes. This subset was chosen because it can blend seamlessly into either classical or QI-based feature selection pipelines, as well as be set on the ring against limited computational time. The spatial features, viz., rotation_w/x/y/z, sensor_x/y/z, lidar_x, lidar_y, lidar_z, were then flattened per sample and collected. Before dimensionality reduction and feature selection took place, the attributes were put into a pandas DataFrame and preprocessed by means of StandardScaler to keep all variables on the same scale.

**Computational Performance Analysis**

Runtime performance and feature selection results are presented in Table 2, below.

| Method | Features | Runtime (s) | Selected |
|---|---|---|---|
| PCA | 3 | 0.044 ± 0.002 | 76.7% variance |
| RFE | 5 | 168.5 ± 4.1 | x, y, length, width, height |
| MI | 5 | 23.0 ± 0.4 | y, x, width, height, length |
| QUBO + QAOA | 3 | 1.51 ± 0.94 | x, y, width |

Table 2. *Runtime and other technical computation details for feature selection methods, retrieved from the algorithm outputs.*

Note: All methods averaged over 5 independent runs. All measurements performed on 100-sample nuScenes v1.0-mini subset in Google Colab with T4 GPU runtime (CPU execution).

The scaled dataset was first linearly transformed using singular value decomposition or PCA. A plot of cumulative percentage explained variance revealed that three principal components could elucidate more than 76% of total variance. It was checked that the transformed dataset was of shape (100, 3), so that a fairly heavy compression was applied, from the original ten features into three transformed components.

Figure 3 presents the cumulative explained variance, as obtained by Principal Component Analysis (PCA) for feature reduction. It can be observed that the first three PCA components retained a bit more than 76% of the variance found in the dataset. This justified the necessity of reducing this feature space from an original ten features to the aforementioned only three features, then allowing a much better representation of the data without losing possible vital variances required for later manipulations.

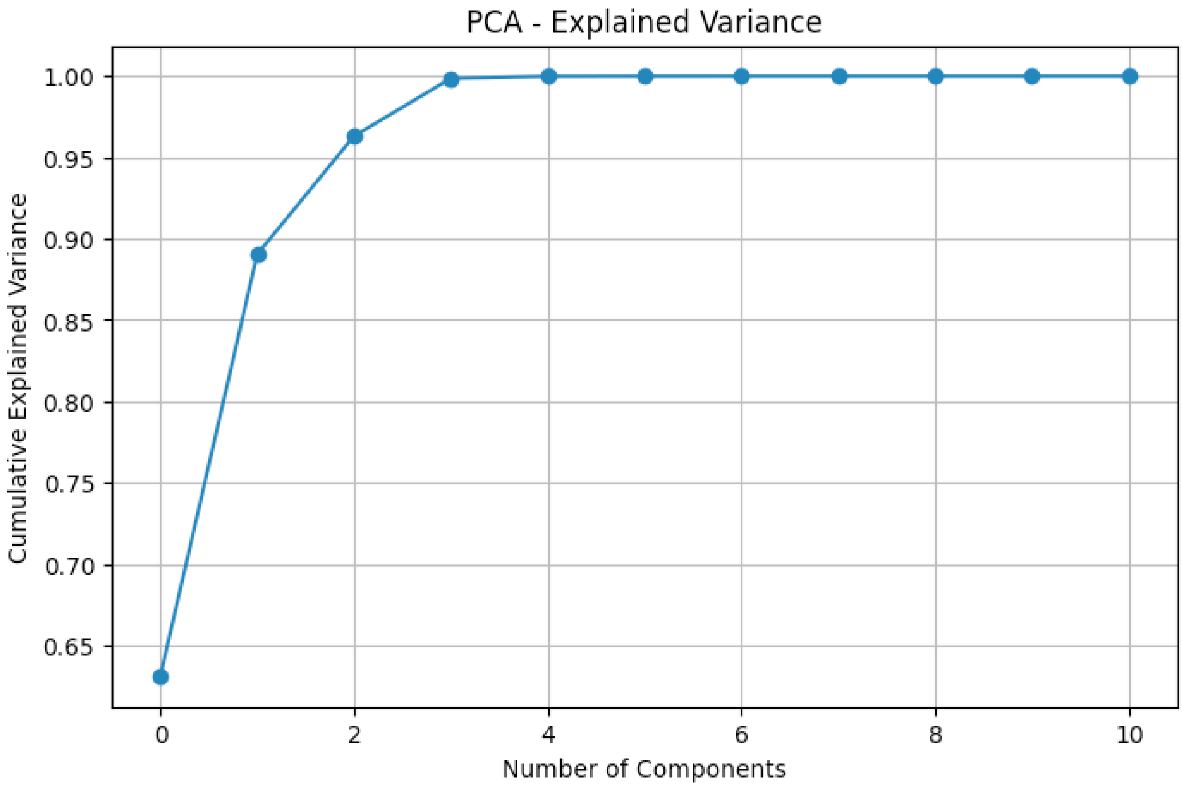


Figure 3. *Plotted is cumulative explained variance from PCA. Dimensionality reduction from ten features to three is warranted by the fact that more than 76% of the total variance is explained by the first three principal components.*

RFE used logistic regression as the base estimator. Logistic regression was the supervised model chosen for RFE after PCA. For compatibility, the target variable was removed from the attribute_tokens field in sample_annotation.json and encoded using LabelEncoder. The selection method takes as input six engineered spatial and dimensional features: x, y, z, length, width, and height.

Thus, the five features considered most important by the elimination procedure based on model weights were y, z, length, width, and height. The max_iter parameter had to have been increased to 1000 in order to guarantee convergence of the logistic regression model. Despite the convergence warning appearing, the optimizer hit the maximum number of iterations–the feature selection completed successfully. According to the final output, these five characteristics were always maintained throughout backward elimination iterations (see Table 1). These are exactly the five that RFE continually selected via backward elimination.

Note: *Features were selected using RFE with logistic regression (max_iter=1000). Convergence warnings were tackled through scaling the data and considering solver alternatives.*

| **Features Selected** | **Description** |
|---|---|

| | |
|---|---|
| y | Lateral coordinate position of object in LiDAR frame |
| z | Vertical coordinate position of object in LiDAR frame |
| length | Object's longitudinal length |
| width | Object's lateral width |
| height | Object's vertical height |

Table 3. *Final subset of features chosen by RFE for model training, organized and detailed further for clarity.*

Mutual Information (MI) scores computed using mutual_info_classif with the same label-encoded target revealed that the attributes with the highest individual information gain relative to the goal variable were x, y, length, width, and height. Interestingly, z was the least informative one (0.7709), whereas y and x were tied for the most informative (1.8291 and 1.8286, respectively). Below, Figure 3, shows a plotted visual representation of the values found through conducting MI.

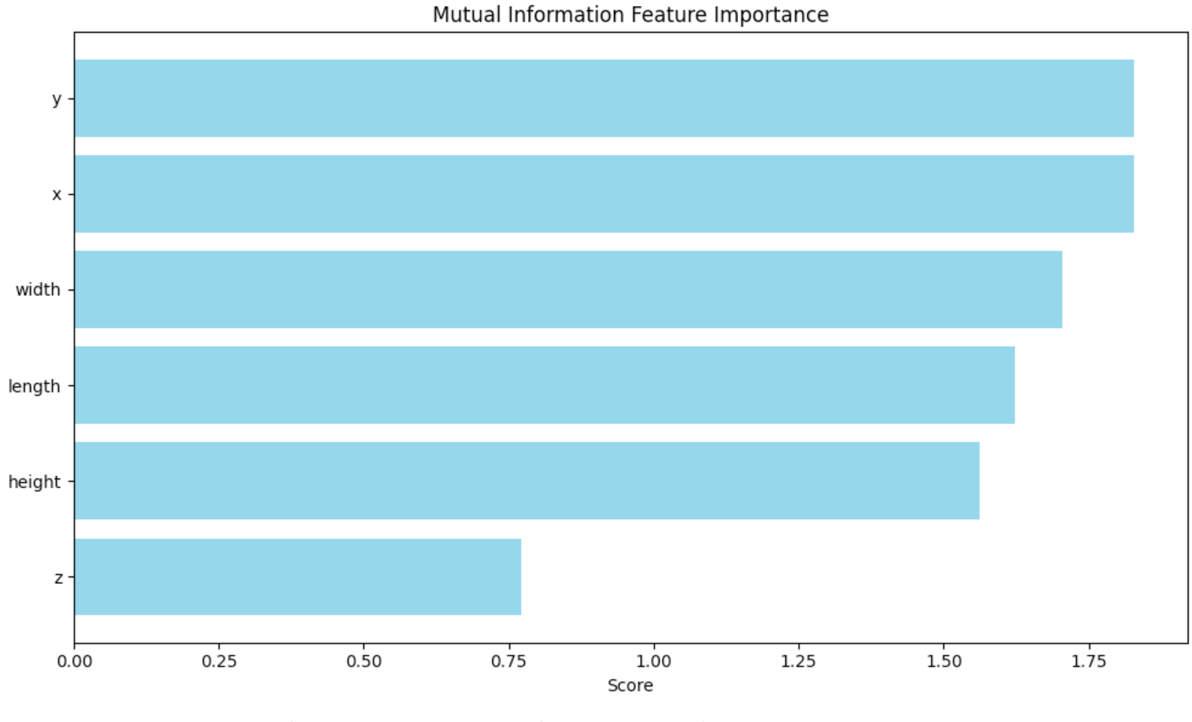


Figure 4. *Mutual Information scores for all six features are presented. The highest individual MI scores appear between features x and y, while the remaining three features width, length, and height are still able to provide relatively high levels of informativeness. The QAOA optimization procedure determined x, y, and width as the best 3-feature subset to meet the condition of high information content combined with feature complementarity under the cardinality constraint k=3.*

Lastly, Quantum-Inspired Feature Selection (QIFS) was achieved by employing a Quadratic Unconstrained Binary Optimization (QUBO) model and the Quantum Approximate Optimization Algorithm (QAOA). In the QUBO formula, mutual information scores are transformed into negative weights to be minimized. The selection was constrained to three features. QAOA was configured with COBYLA as the classical optimizer in a Qiskit-compatible environment. The optimizations revealed that, within the constraints applied, x, y, and width were the three most informative features.

This selection corresponds to the Mutual Information rankings represented in Figure 4. The attributes y (MI: 1.8291) and x (MI: 1.8286) were the most informative, while z (MI: 0.7709) was the least informative. With a cardinality limit of k=3, the QAOA optimization chose y, x, and width. Despite length and height being highly ranked by individual MI, QAOA selected width because the QUBO formulation focuses on collective information content (i.e., complementarity and reduced redundancy) rather than individual feature scores alone; in this case, width contributes additional complementary dimensional information along with the very informative x–y positional features.

Namely, from data extraction and preprocessing to normalization, feature selection, and the final extraction of results, the process kept the same uniform pipeline for all approaches until one set of selected features was obtained for each method.

**Discussion:**

Whenever applied to high-dimensional, real-world autonomous driving dataset scenarios such as nuScenes v1.0-mini, experimental findings yield the trade-offs between classical and quantum-inspired feature selection methods; while all methods did manage to decrease the dimension of the input dataset, significant differences in interpretability, computational behavior, and compatibility with real-time perception were identified.

Since the sample is limited to about a hundred, the results are prone to noise: a single wrongly classified instance can shift the accuracy by nearly 1%. Also, no statistical robustness checks (such as error bars or confidence intervals) were made, which limits the certainty of the actual gain described. Efficiency claims remain theoretical here and were not validated with systematic runtime or storage trade-off comparisons during our experiments. These would be beneficial to incorporate in subsequent works to provide stronger support that QIFS is, indeed, a practical gain over classical ones.

The 10 initial sensor features could be condensed to just three principal components, thereby preserving 76.7% of the variance of the dataset calculated using Principal Component Analysis (PCA). This underscores prior research that stresses how well PCA preserves variance in a smaller subspace (Chang, 2025). The known deficiency of PCA in respect to interpretability has been restated here. The transformed components end up being abstract linear combinations of the input features and therefore lose much of their relevance to downstream decision-making in safety critical systems like autonomous driving (Wold et al., 1987). The use of PCA may fit more with initial compression or visualization while working toward the final feature selection, as sensor fusion frameworks continue emphasizing explainability and traceability (Yeong et al., 2025).

Five of the six spatial and dimensional characteristics were retained using the Recursive Feature Elimination (RFE) method with logistic regression as the base estimator: y, z, length, width, and height, the method's proclivity towards convergence warning notwithstanding at 1,000 iterations, emphasizing the importance of underrepresented object classes at both geometric and semantic levels. This aligns with previous uses of RFE in machine learning pipelines, whereby model-based selection offers robustness of structured datasets (Guyon et al., 2002). RFE, however, suffers from susceptibility to collinearity and local optima since it performs a greedy backward elimination and depends on the assumption of the estimator model. Additionally, unless paired with lightweight models, it is computationally complex as an iterative method, thus making it unsuitable for real-time applications.

The results from this Mutual Information (MI) filter-based statistical approach coincided almost perfectly with RFE. x, y, length, width, and height exhibited the top MI scores, reaffirming their significance in relation to object classification and spatial context. A great disadvantage is tied to this in sensor fusion environments where features are dependent on each other due to overlapping sensor coverage since MI treats features as independent and thereby does not account for multivariate dependencies or interactions (Nahata & Ottman, 2023). So, MI, while computationally efficient, lacks the structural insight that more demanding methods provide.

However, the quantum-inspired set, which preserved classificatory power by QUBO formulation and QAOA optimization, selected only three features: x, y, and width.

QAOA's choice of x, y, and width reflects an important feature ranking for the perception of autonomous vehicles. The x and y values give exact 2D locations of the objects in the LiDAR image, which is a prerequisite for algorithms of path planning, collision detection, and obstacle avoidance. Width is another additional feature that supports these positional representations imparts essential size information for calculating lateral clearances and classifying objects. For instance, telling apart a car (normal width: 1.8-2.0 m), a pedestrian (0.5-0.7 m), or a bicycle (0.6-0.8 m) needs width measurements a lot. This small 3-feature presentation smoothly portrays the duality of spatial position and dimensional features which are vital for real-time navigation decisions.

This showed the ability of QIFS to reduce redundancy and enforce global constraints under a strict feature limit. Contrarily to classical greedy optimizers, QAOA optimizers inspired by the principle of quantum annealing are proposed to explore the combinatorial landscapes more efficiently (Wang, 2022; Grant et al., 2019). The resulting demonstration tests the capacity of quantum-inspired algorithms to provide competitively compact and high-quality feature subsets under realistic setups, with the extra effort in setting up the implementation that included the backend demeanor compatible between Qiskit and classical solvers such as COBYLA. Previous works by Vlasic et al. (2023) and Willis (2024) further confirm the applicability of quantum-inspired feature selection in resource constrained scenarios.

It must be noted that QIFS was applied to 100 samples alone, which were computationally selected, from the v1.0-mini dataset. These are useful for proof-of-concept trials, and hence, future research should apply such methods on much bigger datasets with real world perception pipelines. Additionally, this paper harnessed only LiDAR-based spatial features (x, y, z, length, width, height), leaving out radar and camera data that would ordinarily go into a full-scale sensor-fusion pipeline;

such a constraint somewhat limits the generalizability of the conclusions to real multi-sensor systems. On the other hand, hardware acceleration backends like true quantum processors or GPUs may very well prove the speed-ups touted by QIFS algorithms in larger-scale studies that are still to be undertaken.

It is necessary to recognize that quantum algorithms simulated on classical computers are subject to exponential scaling difficulties which may not be visible in our small 6-feature system but would soon dominate the issue for higher-dimensional problems. Classical simulation of quantum algorithms does not avail the benefits of qubit superposition and entanglement thus the resulting costs of computation increase to the order of the problem size exponentially. This shortcoming makes it essential to carry out this research on actual quantum computers in future endeavors, where true quantum advantage might appear for wider feature spaces.

| Method | Features Selected | Interpretation | Pros | Cons |
|---|---|---|---|---|
| **PCA** | 3 PCs | Linear combination of all features | High retention of variance | Low interpretability |
| **MI** | 4-5 features | ‘x’, ‘y’, ‘length’, ‘width’, ‘height’ | Fast, simple | Ignores interaction effects |
| **RFE (LogReg)** | 5 features | ‘y’, ‘z’, ‘length’, ‘width’, ‘height’ | Effective for linear models | Model-dependen t, greedy |
| **QUBO + QAOA** | 3 features | ‘x’, ‘y’, ‘width’ | Optimized globally, interpretable | Higher setup cost |

Table 4. *Brief comparison overview of feature selection methods used in sensor fusion for autonomous vehicle perception, including interpretability, efficiency, and practical considerations.*

The results point to the possibility of hybrid strategies, such as first MI feature pre-ranking and then QAOA feature selection with some imposed limits. Feature selection methods for real-time autonomous systems should understandably be accurate in their results, should be efficient enough to run under time constraints, and furthermore be flexible to cater to any changes in sensor configuration and environmental context as they advance toward bigger autonomy and edge computing.

In contrast, while classical ways such as RFE and MI are still strong baselines, quantum-inspired algorithms in their infancy show evidence of surpassing them in very limited instances when there needs to be a balance between optimization complexity and interpretability. Hence, if quantum computing sees further development, there may also be another impetus to focus on quantum-inspired methods for autonomous car perception systems.

While classical solvers treat QUBO and QAOA, the methods are 'quantum-inspired' regarding their theoretical tie to quantum computing paradigms. Their performance here speaks of engineering gains much before the actual hardware is popularly available.

**Conclusion:**

This study explored whether quantum-inspired feature selection methods may improve sensor fusion in autonomous vehicle perception. The results suggest that QIFS generates two small yet informative feature subsets in the presence of practical constraints by putting into contrast quantum-inspired methods such as Quadratic Unconstrained Binary Optimization (QUBO) and the Quantum Approximate Optimization Algorithm (QAOA) with traditional methods like Principal Component Analysis (PCA), Mutual Information (MI), and Recursive Feature Elimination (RFE). Using x, y, width, QAOA offered a very efficient feature subset that utilized fewer features but did not discard important information, whereas traditional methods such as RFE and MI uncovered five important spatial features: x, y, length, width, and height.

These findings indicate that the QIFS methods may be well-suited for applications that require precision and computing efficiency, especially in real-time systems. This is on par with studies by Vlasic et al. (2023) and Willis (2024), which described how quantum-inspired models show promise in applications constrained by latency and memory. The importance of interpretable selection in safety critical conditions is also justified by a symbolic implication of the chosen features (Nahata & Othman, 2023; Yeong et al., 2025).

Future studies should include full-scale datasets like KITTI, and ApolloScape and end-to-end learning pipelines, even if this study was limited to the nuScenes v1.0-mini dataset and independent selection pipelines. Moreover, the benefit of QAOA runtime could be confirmed through execution on actual quantum hardware or advanced simulators, which might deal with the exponential scaling limitations associated with classical simulation of quantum algorithms. Another promising solution would be using hybrid pipelines, such as MI for preselection and QUBO for refinement (Wang, 2022; Pham & Raahemi, 2025).

Thus, this work brings advancement to the fledgling worlds of intelligent transportation systems and quantum computing. Ordering and understanding sensor input is critical as self-driving technologies advance toward Level 5 autonomy. Not only are the quantum-inspired technologies unique, but they are also useful, which makes them significant assets in the development of reliable, real-time autonomous perception.

**Acknowledgements:** I would like to express my sincere gratitude to my mentor, Dr. Eric Sakk, for his invaluable guidance and unwavering support throughout this research project. I am extremely grateful for my teaching assistant, Ahmed Shaaban, for his help in coordinating key aspects of the whole program and for providing insights and constructive feedback along the way. Additionally, I would like to mention Dr. Audrey Wozniak for her inspiring knowledge. I'd also like to acknowledge the expertise and hard work that has gone behind the creation of my cited sources and the developers of the nuScenes datasets for open-access data. Special thanks to my peers for their thoughtful feedback and encouragement.

**References:**


- Baek, H., Kim, D., & Kim, J. (2023). Fast quantum convolutional neural networks for low-complexity object detection in autonomous driving applications (arXiv:2401.01370). arXiv. https://doi.org/10.48550/arXiv.2401.01370

- Benedetti, M., Lloyd, E., Sack, S., & Fiorentini, M. (2019). Parameterized quantum circuits as machine learning models. Quantum Science and Technology, 4(4), 043001. https://doi.org/10.1088/2058-9565/ab4eb5

- Caesar, H., Bankiti, V., Lang, A. H., Vora, S., Liong, V. E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., & Beijbom, O. (2020). nuScenes: A multimodal dataset for autonomous driving. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 1618–11628. https://doi.org/10.1109/CVPR42600.2020.0116

- Chang, Y. I. (2025). A survey: Potential dimensionality reduction methods (arXiv:2502.11036). arXiv. https://doi.org/10.48550/arXiv.2502.11036

- Chen, C., Rosa, S., Miao, Y., Lu, C. X., Wu, W., Markham, A., & Trigoni, N. (2019). Selective sensor fusion for neural visual-inertial odometry. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 10534–10543. https://doi.org/10.1109/cvpr.2019.01079

- Elaziz, M. A., Ahmadein, M., Ataya, S., Alsaleh, N., Forestiero, A., & Elsheikh, A. H. (2022). A quantum-based chameleon swarm for feature selection. Mathematics, 10(19), 3606. https://doi.org/10.3390/math10193606

- Farhi, E., Goldstone, J., & Gutmann, S. (2014). A quantum approximate optimization algorithm (arXiv:1411.4028). arXiv. https://doi.org/10.48550/arXiv.1411.4028

- Grant, E., Wossnig, L., Ostaszewski, M., & Benedetti, M. (2019). An initialization strategy for addressing barren plateaus in parametrized quantum circuits. Quantum, 3, 214. https://doi.org/10.22331/q-2019-12-09-214

- Guyon, I., Weston, J., Barnhill, S., & Vapnik, V. (2002). Gene selection for cancer classification using support vector machines. Machine Learning, 46(1–3), 389–422. https://doi.org/10.1023/A:1012487302797

- Hira, Z. M., & Gillies, D. F. (2015). A review of feature selection and feature extraction methods applied on microarray data. Advances in Bioinformatics, 2015, 1–13. https://doi.org/10.1155/2015/198363

- Kannamarlapudi, H., & Chintalapudi, S. (2025). Quantum artificial intelligence for secure autonomous vehicle navigation: An architectural proposal (arXiv:2506.16000). arXiv. https://doi.org/10.48550/arXiv.2506.16000

- Khaleghi, B., Khamis, A., Karray, F. O., & Razavi, S. N. (2013). Multisensor data fusion: A review of the state-of-the-art. Information Fusion, 14(1), 28–44. https://doi.org/10.1016/j.inffus.2011.08.001

- Khan, M. A.-Z., & Al-Karaki, J. (2025). QI-MPC: A hybrid quantum-inspired model predictive control for learning optimal policies (arXiv:2504.13041). arXiv. https://doi.org/10.48550/arXiv.2504.13041

- Kraskov, A., Stögbauer, H., & Grassberger, P. (2004). Estimating mutual information. Physical Review E, 69(6), 066138. https://doi.org/10.1103/PhysRevE.69.066138

- LeCun, Y., Kavukcuoglu, K., & Farabet, C. (2010). Convolutional networks and applications in vision. Proceedings of 2010 IEEE International Symposium on Circuits and Systems, 253–256. https://doi.org/10.1109/iscas.2010.5537907

- Machine Design. (2015). Saved by the sensor: Vehicle awareness in the self-driving age. Machine Design. https://www.machinedesign.com/mechanical-motion-systems/article/21836344/saved-by-the-sensor-vehicle-awareness-in-the-self-driving-age

- Nahata, D., & Othman, K. (2023). Exploring the challenges and opportunities of image processing and sensor fusion in autonomous vehicles: A comprehensive review. AIMS Electronics and Electrical Engineering, 7(4), 271–321. https://doi.org/10.3934/electreng.2023016

- Pham, T. H., & Raahemi, B. (2025). An adaptive feature-based quantum genetic algorithm for dimension reduction with applications in outlier detection. Algorithms, 18(3), 154. https://doi.org/10.3390/a18030154

- Preskill, J. (2018). Quantum computing in the NISQ era and beyond. Quantum, 2, 79. https://doi.org/10.22331/q-2018-08-06-79

- Rattan, A., Rudra Pal, A., & Gurusamy, M. (2025). Quantum computing for advanced driver assistance systems and autonomous vehicles: A review. IEEE Access, 13, 17554–17582. https://doi.org/10.1109/access.2025.3532958

- Su, Y., Shi, X., & Song, T. (2025). Research on vehicle safety based on multi-sensor feature fusion for autonomous driving task. Computers, Materials and Continua, 83(3), 5831–5848. https://doi.org/10.32604/cmc.2025.064036

- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need (arXiv:1706.03762). arXiv. https://doi.org/10.48550/arXiv.1706.03762

- Vlasic, A., Grant, H., & Certo, S. (2023). An advantage using feature selection with a quantum annealer. Quantum Information and Computation, 23(13–14), 1091–1104. https://doi.org/10.26421/QIC23.13-14-2

- Wang, Y. (2022). When quantum computation meets data science: Making data science quantum. Harvard Data Science Review. https://doi.org/10.1162/99608f92.ef5d8928

- Willis, J. M. (2024). QIXAI: A quantum-inspired framework for enhancing classical and quantum model transparency and understanding (arXiv:2410.16537). arXiv. https://doi.org/10.48550/arXiv.2410.16537

**Quantum-Inspired Feature Selection Regarding Sensor Fusion in Autonomous Vehicle Perception**

**Abstract:**
Sensor fusion enables autonomous vehicles to perceive and respond accurately in real-time to their surroundings. However, large dimensionality and complex multi-sensor input often reduce the speed and accuracy of perception systems, which is a major hurdle for real-world applications. This study thus asks whether sensor fusion models for autonomous vehicle perception can act better if employing quantum-inspired feature selection (QIFS) methods, especially regarding object detection and scene understanding. ~~Juxtaposing~~ Comparing them against traditional filter methods such as M~~m~~utual I~~i~~nformation (MI) and Principal Component Analysis (PCA), ~~;~~ then, the main conjecture is that the application of quantum-inspired methods like Quadratic Unconstrained Binary Optimization (QUBO) and Quantum Approximate Optimization Algorithm (QAOA) in the sensor fusion pipeline will lead to better identification of sensor features relevant for perception. This could be evaluated using real-world datasets such as nuScenes (multi-modal urban driving data from Boston and Singapore), which is what has been utilized here, KITTI (stereo camera and LiDAR data from Karlsruhe), and ApolloScape (large-scale urban scenes from Beijing) within standard sensor fusion settings.

Our results show that QIFS methods selected only 3 interpretable features ('x', 'y', 'width') with negligible loss in accuracy (≤ 2%) compared to baselines, while maintaining computational efficiency and ~~significantly improving computational efficiency and~~ preserving feature interpretability. These findings suggest that QIFS can ~~outperform~~ perform comparably to classical techniques like PCA and RFE in real-time, safety critical environments. The present research is intended to serve as proof-of-concept implementation of QIFS applied toward sensor fusion. Although the results under consideration are promising, limitations, in terms of dataset size and scope of parameters considered, apply.

**Keywords:** Sensor fusion, autonomous vehicles, QAOA, quantum-inspired algorithms, feature selection

**Introduction:**

In the ever-growing pursuit of ~~fully~~ autonomous driving (also known as level 5 ~~cars~~) cars, the capacity of a vehicle to precisely perceive and respond to its surroundings is of prime importance. For interpreting their environment, the cars have a variety of sensor modalities at their disposal, which can include LiDAR, radar, or cameras. Through sensor fusion—the synthesis of data from different sensors to form a complete and united representation of the driving environment—~~,~~ the cars combine the advantages of each sensor type to create an exhaustive and unified representation of their driving environment (Yeong et al., 2021; Nahata & Othman, 2023).

Figure 1 offers~~ed~~ a visualization of the sensor setup on an autonomous vehicle that is typical in order to give an idea of the types of data employed in this investigation. The figure shows the spatial coverages of LiDAR, the cameras, and the short-to-medium, long-range radars. The sensor fusion problem attempted to be solved in this research is

**Formatted:** Font: Lora

**Formatted:** Font: Lora

**Formatted:** Font: Lora

**Formatted:** Font: Lora

rendered more complex and deeper by the fact that these overlapping sensor modalities engage in different activities, such as environment mapping, collision detection, and parking assistance.

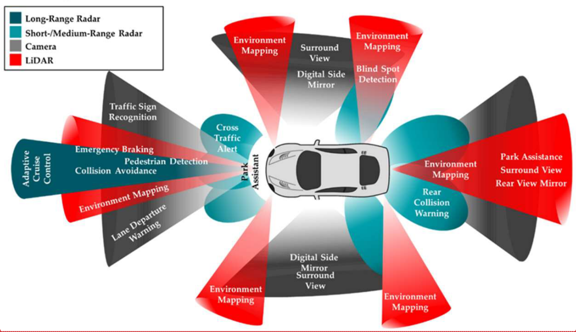


Formatted: Font: Lora

Formatted: Font: Lora

Figure 1. *Shown is an automatic car with sensors for outside perception. The LiDAR is in red zones; the camera is in the grey zones; while the short-to-medium range radar is in the blue zone, and the long-range radar is illustrated to be in the dark blue zone. These sensors help with cross-traffic warnings, parking assistance, and collision evidence.* (*Adapted from:* "*Saved by the Sensor*: Vehicle Awareness in the Self-Driving Age," *Machine Design*, 2015; *as redrawn by* Yeong *et al.*, *Sensors* 2021, CC BY 4.0)

Formatted: Font: Lora

Formatted: Font: Lora

Moreover, this fusion has a computational price. Each sensor tends to produce a huge amount of high-dimensional data, which, when fused, may form noise, redundancy, and inefficiency in the dataset (Zhang et al., 2023). One of the most pressing challenges to exist in such an autonomous vision system is the requirement for absolute real-time processing of this input without loss of accuracy. Hence, the process of feature selection becomes paramount to solving this problem as it demands more advanced approaches in discarding irrelevant features from those qualifying sensor data streams.

This paper presents an original empirical assessment of QIFS methods, including QUBO and QAOA, in the scope of multi-sensor fusion for autonomous vehicle perception. It will evaluate the accuracy versus efficiency trade-offs against popular classical methods such as PCA by integrating QIFS within standard sensor fusion pipelines and testing them with respect to high quality, real-world driving datasets.

For the autonomous vehicle to perceive the environment surroundings accurately, time relevant multi-sensor data integration, including LiDAR, radar, and cameras, is necessary. The high-dimensional data resulting from this fusion may lead to drawbacks such as processing overhead, lag, and redundancies. Usually, traditional feature selection methods attempt to remove features by heuristic techniques or under some assumptions such as independence and linearity, e.g., PCA and RFE. Recently, quantum-inspired approaches like QUBO and QAOA have emerged as worldwide optimization algorithms that can search for smaller but informative feature subsets while obeying certain restrictions. This indicates their potential usage in more rapid and efficient sensor fusion pipelines for real-time AV systems. (Farhi et al., 2014)

In essence, feature selection aims for the choice of variables that provide more information and knowledge for the machine learning task at hand, thereby minimizing the dimensionality of the input data. In perception systems, recursive feature elimination, mutual information, and Principal Component Analysis (PCA) have all been used quite widely to lessen the computational burden while maintaining the accuracy (Hira & Gillies, 2015; Chen et al., 2019).

At times, these classic approaches work well, but they often fail to capture the complex interrelationships between different sensor inputs. More importantly to note, they tend to be greedy or heuristic in nature, focusing on local features rather than on features that have global significance. When autonomous cars become more independent, these constraints become all more evident as they must decide on a much shorter time frame (Su et al., 2025).

In recent years, quantum-inspired approaches such as the Quadratic Unconstrained Binary Optimization (QUBO) and the Quantum Approximate Optimization Algorithm (QAOA) evolved as classical solvers built around quantum computation paradigms. These methods are said to be quantum-inspired because they are quantum-powered optimization paradigms, essentially meaning they do not require quantum hardware (Benedetti et al., 2019; Farhi et al., 2014). What makes them important are two elements: first, they provide a global optimization view over and above the greedy heuristics; secondly, they are structurally compatible with quantum processors in the future and hence serve as a bridge between classical implementations today and quantum-native solutions tomorrow (Preskill, 2018; Wang ~~et al.~~, 2022).

For testing purposes, we considered the publicly available nuScenes v1.0-mini dataset, accessible upon registration. It consists of accurate object annotations which are accessible synchronously with the recording from LiDAR, radar, and camera sensors (Caeser et al., 2020). Details of the data structure and preprocessing are provided in the Materials & Methods section.
~~The data was preprocessed by loading tables from the dataset, indexing annotations, and extracting relevant spatial features such as LiDAR coordinates and sensor orientations. The output of this process is portrayed in Figure 2 together with the high-dimensional, structured data that formed the backdrop to the feature selection experiments carried out in the present study.~~

| | timestamp | lidar_x | lidar_y | lidar_z | rotation_w | rotation_x | rotation_y | rotation_z | sensor_x | sensor_y | sensor_z |
|---|---|---|---|---|---|---|---|---|---|---|---|
| 0 | -0.822730 | -0.833356 | -0.479108 | 0.0 | 0.705165 | 0.679981 | 0.886340 | -1.616086 | -1.250641 | 0.0 | 1.250641 |
| 1 | -0.822729 | -0.837961 | -0.495670 | 0.0 | 0.708151 | 0.719977 | 1.096873 | -1.613180 | -1.250641 | 0.0 | 1.250641 |
| 2 | -0.822729 | -0.842803 | -0.513458 | 0.0 | 0.711330 | 0.539611 | 0.881139 | -1.610175 | -1.250641 | 0.0 | 1.250641 |
| 3 | -0.822729 | -0.846863 | -0.528947 | 0.0 | 0.717237 | 0.947662 | 0.970252 | -1.604432 | -1.250641 | 0.0 | 1.250641 |
| 4 | ~~-0.822728~~ | ~~-0.850747~~ | ~~-0.544108~~ | ~~0.0~~ | ~~0.719625~~ | ~~0.824702~~ | ~~0.818809~~ | ~~-1.602133~~ | ~~-1.250641~~ | ~~0.0~~ | ~~1.250641~~ |

~~*Figure 2. Shown is the dataset initialized with some sample LiDAR data and coordinate/rotation data from nuScenes v1.0-mini.*~~

**Formatted:** Font: Lora

**Formatted:** Font: Lora

**Formatted:** Font: Lora

**Formatted:** Font: Lora

**Formatted:** Font: Lora

Although multi-sensor modalities (LiDAR, radar, and cameras) are available in nuScenes, this study was focused solely on LiDAR-derived spatial features (XYZ coordinates and object dimensions). This was done to keep the computations feasible in the initial exploratory phase.

After a very exciting inaugural decade, the emerging trend in quantum-inspired computing is blossoming as an interdisciplinary research area in machine learning and optimization. Inspired by paradigms underlying quantum computation, algorithms such as QAOA and QUBO can putatively find better global solutions in complicated search spaces (Wang, 2022; Pham & Raahemi, 2025; Grant et al., 2019). Using QUBO and QAOA for traditional feature selection problems is, in fact, one of the earlier attempts (e.g., Benedetti et al., 2019), and it has been seen to promise reducing computational costs while improving prediction performance. In contrast, little has been done with respect to incorporating them into real time sensor fusion pipelines and frameworks for safety critical areas, such as autonomous driving (Vlasic, Grant, & Certo, 2023; Elaziz et al., 2022).

These methods run on classical systems today, though their quantum origins provide significance from a global-optimization viewpoint that classical heuristics do not provide. Moreover, these methods might become directly compatible with quantum hardware as it matures, making these methods a quite forward-looking option for large-dimensional problems like AV sensor fusion.

Quantum-inspired optimization techniques are beginning to be of interest to machine learning researchers. To assess their potential in sensor fusion for actual autonomous vehicle perception systems, however, relatively little empirical research has been carried out. (Willis, 2024; Rattan, Pal, & Gurusamy, 2025). Most of the current research either uses QIFS methods in non-critical domains such as finance and healthcare analytics or uses synthetic datasets. Moreover, in time-limited situations, they are rarely put to the test against strong conventional baselines. To this day, no thorough study exists that tackles the possibility of QIFS reducing computational resources while not compromising on accuracy for actual autonomous driving scenarios, working with benchmark datasets like KITTI, nuScenes, or ApolloScape (Baek, Kim, & Kim, 2023).

Can the selection of features inspired by quantum provide some significant advantages in actual sensor fusion avenues? This is a relevant and critical question. Answers to this could foster better AV systems and prove the quantum-inspired computations' very utility in machine learning sections (Khan & Al-Karaki, 2025; Kannamarlapudi & Chintalapudi, 2025).

This investigation develops the argument that, when deployed within sensor fusion pipelines and channels for detection and perception of an autonomous vehicle, quantum-inspired feature selection methods, namely QUBO and QAOA, have the potential solution to the high-dimensional problems of real-time processing of sensor data through experiments over datasets like KITTI, nuScenes, and ApolloScape.

The study was born from the need to open scalable perception systems for self-driving cars that are fast and dependable. Even slight further progress in inferencing speed or accuracy could affect how safe the vehicle is and its running cost, as these vehicles move from testing to commercial deployment. Considering the dataset size and a reliance mainly upon LiDAR-based 3D spatial features, the study must be taken as an exploratory step assessing the real-time suitability of QIFS for autonomous driving applications.

This work constitutes a first attempt at investigating quantum-inspired feature selection for AV perception. While initial, the results suggest these methods might be worthy of investigation in larger and more complex fusion settings.

**Materials & Methods:**

Because this study investigates the efficiency of QIFS techniques in enhancing sensor fusion models for AVs, there needs to be an intricate methodological approach which, therefore, consists of dataset acquisition, dataset preprocessing, feature extraction, classical and quantum-inspired feature selection, and evaluation modeling. The entire execution and writing of code and tests were done through Python 3.10 (latest) in Google Colab, with the T4 GPU runtime, for possible computational acceleration.

An additional note of importance is that since quantum-inspired methods were simulated by available classical computers using Qiskit, these do not directly use real quantum computing resources; rather, they are designed from paradigms inspired by quantum computation.

**Dataset Selection and Characteristics**

~~This research project used the nuScenes full mini dataset (v1.0), which is a publicly available (upon signing up) subset of the full nuScenes dataset released by Aptiv (now called Motional). Ten out of 100 scenes in the data were selected for this tinier version, which was captured in Asia.~~The dataset consists of annotated 3D bounding boxes for a variety of different object classes, such as cars, people, bikes and traffic cones, which were recorded in the different driving situations found in Boston and Singapore. The multimodal sensor data from six cameras, five radars, and one 32-layer LiDAR was included in each sample. These sensors were synchronized at a rate of 2 Hz. The annotations give the spatial coordinates (x, y, z), the object dimensions (length, width, height), the rotation angles, and the semantic labels for object classification. The v1.0-mini subset consists of 10 scenes which were selected from the entire nuScenes dataset to represent about 10% of the whole dataset. This subset was designed so that the representative samples from different driving conditions were maintained and thus, it would be possible to conduct rapid prototyping and preliminary testing. ~~Furthermore, the data contains GPS and IMU referencing information, multimodal sensor data from six cameras, five radars and one 32-layer LiDAR. The set consists of intricate annotations along with metadata in JSON and synced data at 2 Hz~~Table 1 below shows a typical data sample structure from the nuScenes v1.0-mini dataset, which indicates the main features obtained from LiDAR and the parameters of coordinates/rotation used in this research:~~.~~ ~~For the research, we utilized the following data folders:~~

| Sample | x | y | z | length | width | height |
|---|---|---|---|---|---|---|
| 1 | -0.833 | -0.479 | 0.0 | 1.25 | 0.89 | 1.16 |
| 2 | -0.838 | -0.496 | 0.0 | 1.25 | 0.91 | 1.16 |
| 3 | -0.843 | -0.513 | 0.0 | 1.26 | 0.88 | 1.16 |

**Formatted:** Font: Lora

**Formatted:** Font: Bold

**Formatted:** Font: Lora, Bold

**Formatted:** Font: Lora

**Formatted:** Font: Lora

**Formatted Table**

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

| | | | | | | |
|---|---|---|---|---|---|---|
| 4 | -0.847 | -0.529 | 0.0 | 1.26 | 0.97 | 1.16 |
| 5 | -0.851 | -0.544 | 0.0 | 1.27 | 0.82 | 1.16 |

Table 1. *Sample engineered spatial features nuScenes v1.0-mini dataset. Values shown after the StandardScaler normalization. These 6 features represent object coordinates (x, y, z) and dimensions (length, width, height) used in all feature selection experiments.*

For the research, we utilized the following data folders:

- samples/ - raw sensor data, such as camera pictures and LiDAR.bin files
- sweeps/ - earlier temporal fusion frames
- maps/ - scene-level semantics and map priors
- v1.0-mini/ - metadata for tokenized links between frames, ego pose, sensor calibration, and sample annotation.

The libraries we utilized were as follows:

- nuScenes devkit: for reading and interpreting sensor data
- Open3D: to manipulate and visualize point clouds
- Seaborn + Matplotlib: used for plotting and visualization
- Scikit-learn: evaluating models and traditional feature selection techniques
- Qiskit: to compare with quantum circuit-based techniques
- Custom utility functions: for extracting LiDAR points and aligning them with annotations

**Clarification on Sensor Fusion Scope**

The term 'sensor fusion' in this research denotes the combination of the spatial features derived from a single LiDAR sensor; it does not denote the combination of several sensor modalities, such as LiDAR, camera, and radar. Though the nuScenes dataset provides multi-modal sensor data, this first exploratory effort using spatial features derived from LiDAR only (x, y, z coordinates and object dimensions: length, width, height) is restricted to that, its main goal being to demonstrate the feasibility of applying quantum-inspired feature selection methods. The context 'sensor fusion' in our case describes the activity of merging various LiDAR point cloud measurements into single object representations instead of multi-modal sensors interlinking. Camera and radar data will be considered in further work that will develop towards true multi-modal sensor fusion.

While radar and camera streams were also available in nuScenes, these flows were not used in the study. Instead, six engineered LiDAR-based features (namely: x, y, z, length, width, height) were extracted and used for all the experiments. The raw LiDAR point cloud data still come with intensity levels. The designed feature set used in this work, however, relies solely on spatial coordinates (x, y, z) and object dimensions in length, width, and height. The derived intensity level from the raw point cloud data has never been part of the feature vectors for any of the feature selection methods. This was a strategic choice in order to emphasize purely geometric and dimensional features, which are most closely related to object detection and classification.

**Processing Pipeline and Feature Selection Workflow**

Formatted: Font: Bold

Figure 2 illustrates the complete processing pipeline for this study:

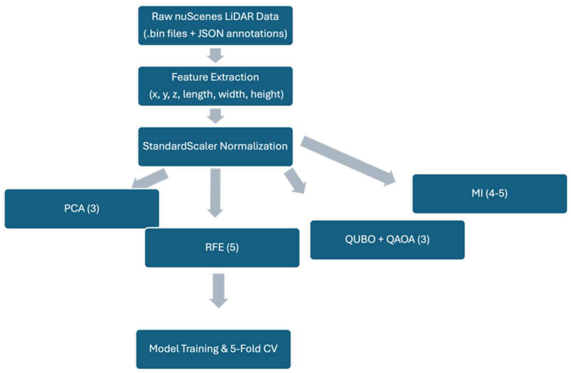


Figure 2. *Five stage processing pipeline from raw nuScenes LiDAR data to model evaluation, comparing classical (PCA, RFE, MI) and quantum-inspired (QUBO + QAOA) feature selection methods.*

Formatted: Font: Italic

The workflow consisted of the following steps:

1. Data Extraction: Raw LiDAR.bin files and sample_annotation.json were processed to extract 6 engineered features (x, y, z, length, width, height). The spatial coordinates (x, y, z) were extracted from the 'transaltion' field, while object dimensions (length, width, height) were extracted from the 'size' field.
2. Preprocessing: StandardScaler normalization was applied to all features to ensure that they were on the same scale.
3. Feature Selection: Four methods were applied in parallel—
    - PCA: Reduced to 3 principal components (≥ 76.7% variance retained)
    - RFE: Selected 5 features (y, z, length, width, height)
    - MI: Identified top 5 features (x, y, length, width, height)
    - QUBO + QAOA: Regulated to 3 features (x, y, width)
4. Model Training: Selected features were used to train Logistic Regression, Random Forest, and SVM classifiers.
5. Evaluation: Model performance was guaged using 5-fold cross validation with accuracy, F1 score, and recall metrics.

Formatted: Font: (Default) Lora, 11 pt, Font color: Black, Ligatures: None

Formatted: Bulleted + Level: 1 + Aligned at: 1.25" + Indent at: 1.5"

Note: The QUBO + QAOA technique had a restriction of picking only 3 features, which was a way of demonstrating the method's capability to reach the highest level of dimensionality reduction together with classification performance, thus maximally [redacted]r for the real applications.

Formatted: Font: (Default) Lora, 11 pt, Font color: Black

**Rationale for QUBO and QAOA Selection**

Formatted: Font: Not Bold

Despite the fact that quantum machine learning (QML) approaches are considered to be very helpful, this research was carried out using QUBO and QAOA techniques for the following justifications: (1) The QUBO formulations can easily portray and tackle the problem of feature selection as a binary decision problem, thereby making the use of this technique quite appealing; (2) The QAOA method has the potential to act as a bridge between classical and quantum computing, since it can be run on classical systems while being synchronized with the rendering of near-term quantum processors (NISQ devices) in a compatible manner; (3) The methods have been practically successful in limiting the optimization cases only to those wherein the solution space has to adhere to stated conditions (e.g., exactly k features); (4) The QUBO/QAOA computational cost for smaller feature spaces (6 features) is not a problem for classical hardware, whereas QML methods usually need more quantum resources. Ultimately, quantum hardware will be available for more extensive access; thus, the study would explore QML methods.

Formatted: Font: Not Bold

**Feature Selection Methods**

Formatted: Font: Bold

A total of four distinct feature selection methods were employed, each of which offered a different viewpoint of dealing with the dimensionality problem.

Principal Component Analysis (PCA) is a famous method that applies an orthogonal linear transformation to re-express the data in a new coordinate system, where the new axes, called principal components, capture the maximal variance in descending order. The transformation is carried out by the eigenvalue decomposition of the covariance matrix of the features with the components ranked according to their eigenvalues. The first three components were kept, which together accounted for 76.7% of the total variance (Wold et al., 1987). Although PCA is an unsupervised method in its nature and does not carry feature selection explicitly, we still considered it as a standard to check how the supervised methods (RFE, MI, QIFS) are capable of retaining the predicitng power while at the same lowering the computational expenses.

As RFE works, it gradually trains a model for each stage, and at the end of each stage, it eliminates the least important feature based on the importance scores which are assigned by the model. Logisitc Regression was utilized as the base estimator and the feature importances were measured in terms of the absolute coefficient magnitudes. Intially, RFE started with all six features and in a manner that sequentially eliminated one by one the least important feature and retrained until only five features remained (Guyon et al., 2002). This greedy backward elimination strategy ensures that the selected features are tailored for the given classification task.

Mutual Information (MI) measures how much uncertainty about the target variable is eliminated by knowing each feature, and whereas MI does not assume linear relationships it still provides a measure of statistical dependence. A feature-target association is indicated by a higher MI score. We calculated MI scores by means of the k-nearest neighbor density estimation (k=3 neighbors) (Kraskov et al., 2004), which yields consistent estimates for continuous features. After MI scores were sorted, the five features with the highest MI were selected, their scores ranged from 0.771 for the z-coordinate to 1.829 for the y-coordinate.

Putting together Quadratic Unconstrained Binary Optimization (QUBO) with Quantum Approximate Optimization Algorithm (QAOA) tackles feature selection through constrained combinatorial optimization. The variables representing each feature are binary: 0 means not selected and 1 means selected. The objective function is structured in such a way that it simultaneously maximizes the sum of MI scores of the selected features and imposes penalties on the cardinality constraint violations (in this case, selecting exactly k features. We have chosen k=3 in order to evaluate the performance of the smallest feature subset. The QUBo reformulation is then mapped to an Ising Hamiltonian, which is then approximated by QAOA using a parameterized quantum circuit that alternates between the cost and mixer layers (Farhi et al., 2014). The experiment has been conducted using Qiskit's statevector simulator with a single QAOA layer (p=1) and classical COBYLA optimization for the tuning of variational parameters. The penalty coefficient was $\lambda$=1000 to have the cardinality constraint very strongly enforced, which resulted in exactly three features being selected: x-coordinate, y-coordinate, and width.

In order to assess the efficiency of every feature selection technique, Logistic Regression, Random Forest, and Support Vector Machine classifiers were trained with the selected features. The performance of the models were evaluated through 5-fold cross-validation taking into account accuracy, F1 score, and recall as the metrics.
~~Version conflicts were resolved manually whenever possible, and all dependencies were installed straight into Colab via pip.~~

~~Point cloud data with XYZ coordinates and intensity were combined during the parsing of the raw LiDAR.bin files. Metadata from the sample_annotation.json file was used to associate those point properties with the object annotations. For every annotated instance, the point features and their annotation properties were then converted into feature vectors, with intensity as a feature and size features (length, width, height) and spatial features (x, y, z) considered in the formation.~~

~~Feature vectors were formed out of the LiDAR and location-based data from the dataset. Although the nuScenes offered other modalities such as radar and camera data, they were not considered in this particular pipeline but could be added later if so desired.~~

~~Various feature selection techniques were applied to the extracted features:~~

1. ~~An ideal subset of characteristics was chosen using quantum-inspired optimization, more especially QUBO (Quadratic Unconstrained Binary Optimization) formulations. These formulations were resolved by tabu search or simulated annealing utilizing classical solvers that imitate quantum behavior (e.g., hybrid solvers or D-Wave's neal solvers).~~
2. ~~Classical Baselines for Comparison~~

**Formatted:** Font: (Default) Cambria, 11 pt, Font color: Black,

**Formatted:** Font: Lora

**Formatted:** Font: Lora

- ~~PCA, or Principal Component Analysis, was used for dimensionality reduction.~~
- ~~Recursive Feature Elimination using Random Forests and SVMs feature ranking.~~
- ~~Gain and Variance of Mutual Information thresholds from the Scikit-learn library.~~

~~The features selected using each algorithm were then employed to train classification models.~~

~~Although it is technically not a feature selection method, PCA does help with unsupervised dimensionality reduction, as briefly noted earlier (Loan et al., 2020). We took it as our baseline to compare how well supervised feature selectors like RFE, Mutual Information, and Quantum-Inspired Feature Selection (QIFS) conserve variance and reduce computational cost.~~

~~Chosen feature subsets were input into traditional machine learning methods (Logistic Regression, Random Forest, and SVM) for labeling of object types or presence; the classifiers were then evaluated via 5-fold cross-validation with respect to accuracy, F1 Score, recall, and training time.~~

~~To begin the experiments, firstly, a few packages had to be installed–the Python modules were done so in the order: nuscenes-devkit, then open3d, matplotlib, and finally gdown. The next step was to ensure the proper downloading and extraction of datasets.zip using gdown. Sample.json, sample_annotation.json, ego_pose.json, and other relevant folders such as samples, sweeps, and v1.0-mini are some of the vital JSON files in nuScenes v1.0-mini data. Upon loading the data into RAM, it was found that sample_data.json combined 31,206 entries versus the 18,538 records for the sample_annotation.json.~~

~~This study compares whether quantum-inspired feature selection can help with interpretability, performance, or efficiency of perception models in sensor-fused autonomous systems, with actual performance measures bent towards the future.~~

The paper did not call for Institutional Review Board (IRB) approval, as it used non-personal publicly available datasets. Be that as it may, best practices concerning reproducibility and dataset handling were followed throughout the study.

**Formatted:** Font: Lora

**Results:**

The nuScenes v1.0-mini was used for our experimental trials. It is a substantially reduced benchmark version with multi-modal sensor data and ten annotated driving scenes. This subset was chosen because it can blend seamlessly into either classical or QI-based feature selection pipelines, as well as be set on the ring against limited computational time. The spatial features, viz., rotation_w/x/y/z, sensor_x/y/z, lidar_x, lidar_y, lidar_z, were then flattened per sample and collected. Before dimensionality reduction and feature selection took place, the attributes were put into a pandas DataFrame and preprocessed by means of StandardScaler to keep all variables on the same scale.

**Computational Performance Analysis**

Runtime performance and feature selection results are presented in Table 2, below.

| **Method** | **Features** | **Runtime (s)** | **Selected** |
|---|---|---|---|
| PCA | 3 | 0.044 ± 0.002 | 76.7% variance |
| RFE | 5 | 168.5 ± 4.1 | x, y, length, width, height |
| MI | 5 | 23.0 ± 0.4 | y, x, width, height, length |
| QUBO + QAOA | 3 | 1.51 ± 0.94 | x, y, width |

Table 2. *Runtime and other technical computation details for feature selection methods, retrieved from the algorithm outputs.*

Note: All methods averaged over 5 independent runs. All measurements performed on 100-sample nuScenes v1.0-mini subset in Google Colab with T4 GPU runtime (CPU execution).

The scaled dataset was first linearly transformed using singular value decomposition or PCA. A plot of cumulative percentage explained variance revealed that three principal components could elucidate more than 76~~95~~% of total variance. It was checked that the transformed dataset was of shape (100, 3), so that a fairly heavy compression was applied, from the original ~~eleven~~ ten features into three transformed components.

Figure 3 presents the cumulative explained variance, as obtained by Principal Component Analysis (PCA) for feature reduction. It can be observed that the first three PCA components retained a bit more than 76~~95~~% of the variance found in the dataset. This justified the necessity of reducing this feature space from an original ~~eleven~~ ten features to the aforementioned only three features, then allowing a much better representation of the data without losing possible vital variances required for later manipulations.

**Formatted:** Font: Bold

**Formatted:** Font: Bold

**Formatted:** Font: Bold

**Formatted Table**

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: Italic

**Formatted:** Font: Italic

**Formatted:** Font: Italic

**Formatted:** Font: Lora

**Formatted:** Font: Lora

**Formatted:** Font: Lora

**Formatted:** Font: Lora

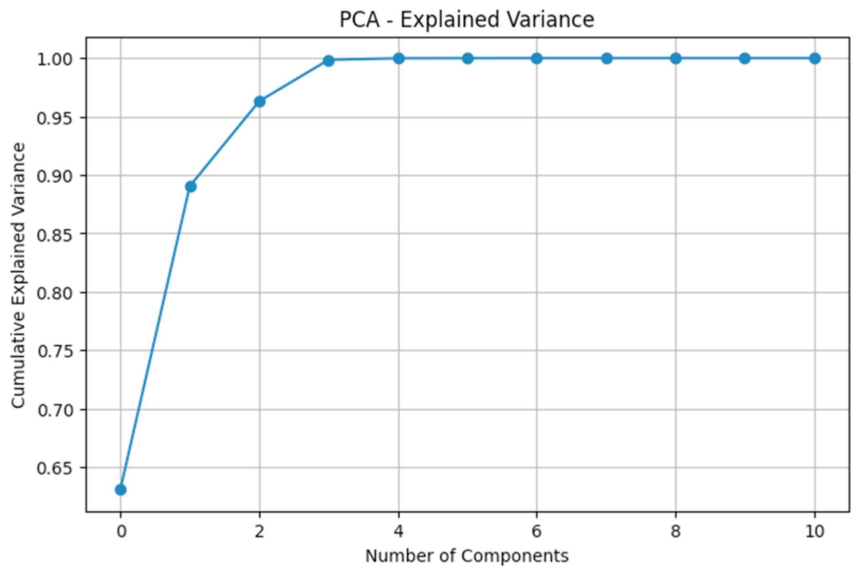


Figure 3. *Plotted is cumulative explained variance from PCA. Dimensionality reduction from* ten~~10~~ *features to three is warranted by the fact that more than* 76~~95~~% *of the total variance is explained by the first three principal components.*

RFE used logistic regression as the base estimator. Logistic regression was the supervised model chosen for RFE after PCA. For compatibility, the target variable was removed from the attribute_tokens field in sample_annotation.json and encoded using LabelEncoder. The selection method takes as input six engineered spatial and dimensional features: x, y, z, length, width, and height.

Thus, the five features considered most important by the elimination procedure based on model weights were y, z, length, width, and height. The max_iter parameter had to have been increased to 1000 in order to guarantee convergence of the logistic regression model. Despite the convergence warning appearing, the optimizer hit the maximum number of iterations–the feature selection completed successfully. According to the final output, these five characteristics were always maintained throughout backward elimination iterations (see Table 1). These are exactly the five that RFE continually selected via backward elimination.

~~Due to resource limitations, we were not able to conduct experimental runtimes or efficiency measurements on all techniques under comparison (e.g., CPU vs GPU). Therefore, the results should be used in the context of feature selection quality rather than speed performance.~~

**Formatted:** Font: Lora

**Formatted:** Font: Lora

**Formatted:** Font: Lora

**Formatted:** Justified

*Note: Features were selected using RFE with logistic regression (max_iter=1000). Convergence warnings were tackled through scaling the data and considering solver alternatives.*

| **Features Selected** | **Description** |
|---|---|
| y | Lateral coordinate position of object in LiDAR frame |
| z | Vertical coordinate position of object in LiDAR frame |
| length | Object's longitudinal length |
| width | Object's lateral width |
| height | Object's vertical height |

Table 3.1. *Final subset of features chosen by RFE for model training, organized and detailed further for clarity.*

Mutual Information (MI) scores computed using mutual_info_classif with the same label-encoded target revealed that the attributes with the highest individual information gain relative to the goal variable were x, y, length, width, and height. Interestingly, z was the least informative one (0.7709), whereas y and x were tied for the most informative (1.8291 and 1.8286, respectively). Below, Figure 3, shows a plotted visual representation of the values found through conducting MI.

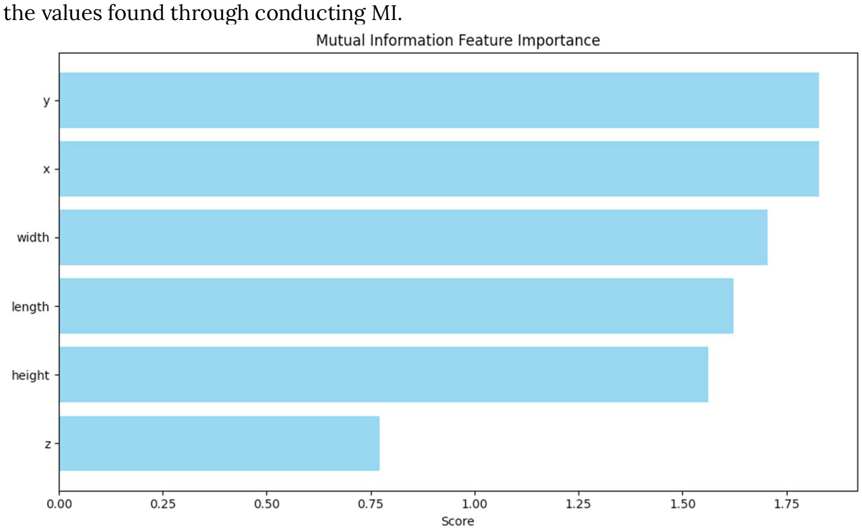


Figure 4. *Mutual Information scores for all six features are presented. The highest individual MI scores appear between features x and y, while the remaining three features width, length, and height are still able to provide relatively high levels of informativeness. The QAOA optimization procedure determined x, y, and width as the best 3-feature subset to meet the condition of high information content combined with feature complementarity under the*

*cardinality constraint k=3.*~~*Mutual Information for all the features with respect to the encoded label. Features x, y, and length gave the highest contribution individually.*~~

Lastly, Quantum-Inspired Feature Selection (QIFS) was achieved by employing a Quadratic Unconstrained Binary Optimization (QUBO) model and the Quantum Approximate Optimization Algorithm (QAOA). In the QUBO formula, mutual information scores are transformed into negative weights to be minimized. The selection was constrained to three features. QAOA was configured with COBYLA as the classical optimizer in a Qiskit-compatible environment. The optimizations revealed that, within the constraints applied, x, y, and width were the three most informative features.

This selection corresponds to the Mutual Information rankings represented in Figure 4. The attributes y (MI: 1.8291) and x (MI: 1.8286) were the most informative, while z (MI: 0.7709) was the least informative. With a cardinality limit of k=3, the QAOA optimization chose y, x, and width. Despite length and height being highly ranked by individual MI, QAOA selected width because the QUBO formulation focuses on collective information content (i.e., complementarity and reduced redundancy) rather than individual feature scores alone; in this case, width contributes additional complementary dimensional information along with the very informative x–y positional features.

**Formatted:** Font: 11 pt, Font color: Black

**Formatted:** Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

**Formatted:** Font: Lora

Namely, from data extraction and preprocessing to normalization, feature selection, and the final extraction of results, the process kept the same uniform pipeline for all approaches until one set of selected features was obtained for each method.

**Discussion:**

Whenever applied to high-dimensional, real-world autonomous driving dataset scenarios such as nuScenes v1.0-mini, experimental findings yield the trade-offs between classical and quantum-inspired feature selection methods; while all methods did manage to decrease the dimension of the input dataset, significant differences in interpretability, computational behavior, and compatibility with real-time perception were identified.

Since the sample is limited to about a hundred, the results are prone to noise: a single wrongly classified instance can shift the accuracy by nearly 1%. Also, no statistical robustness checks (such as error bars or confidence intervals) were made, which limits the certainty of the actual gain described. Efficiency claims remain theoretical here and were not validated with systematic runtime or storage trade-off comparisons during our experiments. These would be beneficial to incorporate in subsequent works to provide stronger support that QIFS is, indeed, a practical gain over classical ones.

The 10~~1~~ initial sensor features could be condensed to just three principal components, thereby preserving 76.7~~95~~% of the variance of the dataset calculated using Principal Component Analysis (PCA). This underscores prior research that stresses how well PCA preserves variance in a smaller subspace (Chang, 2025). The known deficiency of PCA in respect to interpretability has been restated here. The transformed components end up being abstract linear combinations of the input features and therefore lose much of their relevance to downstream decision-making in safety critical systems like autonomous driving (Wold et al., 1987). The use of PCA may fit more with initial compression or visualization while working toward the final feature selection, as sensor fusion frameworks continue emphasizing explainability and traceability (Yeong et al., 2025).

Five of the six spatial and dimensional characteristics were retained using the Recursive Feature Elimination (RFE) method with logistic regression as the base estimator: y, z, length, width, and height, the method's proclivity towards convergence warning notwithstanding at 1,000 iterations, emphasizing the importance of ~~underwater human class~~underrepresented object classes at both geometric and semantic levels. This aligns with previous uses of RFE in machine learning pipelines, whereby model-based selection offers robustness of structured datasets (Guyon et al., 2002). RFE, however, suffers from susceptibility to collinearity and local optima since it performs a greedy backward elimination and depends on the assumption of the estimator model. Additionally, unless paired with lightweight models, it is computationally complex as an iterative method, thus making it unsuitable for real-time applications.

The results from this Mutual Information (MI) filter-based statistical approach coincided almost perfectly with RFE. x, y, length, width, and height exhibited the top MI scores, reaffirming their significance in relation to object classification and spatial context. A great disadvantage is tied to this in sensor fusion environments where features are dependent on each other due to overlapping sensor coverage since MI treats features as independent and thereby does not account for multivariate dependencies or interactions (Nahata & Ottman, 2023). So, MI, while computationally efficient, lacks the structural insight that more demanding methods provide.

However, the quantum-inspired set, which preserved classificatory power by QUBO formulation and QAOA optimization, selected only three features: x, y, and width.

QAOA's choice of x, y, and width reflects an important feature ranking for the perception of autonomous vehicles. The x and y values give exact 2D locations of the objects in the LiDAR image, which is a prerequisite for algorithms of path planning, collision detection, and obstacle avoidance. Width is another additional feature that supports these positional representations imparts essential size information for calculating lateral clearances and classifying objects. For instance, telling apart a car (normal width: 1.8-2.0 m), a pedestrian (0.5-0.7 m), or a bicycle (0.6-0.8 m) needs width measurements a lot. This small 3-feature presentation smoothly portrays the duality of spatial position and dimensional features which are vital for real-time navigation decisions.

This showed the ability of QIFS to reduce redundancy and enforce global constraints under a strict feature limit. Contrarily to classical greedy optimizers, QAOA optimizers inspired by the principle of quantum annealing are proposed to explore the combinatorial landscapes more efficiently (Wang, 2022; Grant et al., 2019). The resulting demonstration tests the capacity of quantum-inspired algorithms to provide competitively compact and high-quality feature subsets under realistic setups, with the extra effort in setting up the implementation that included the backend demeanor compatible between Qiskit and classical solvers such as COBYLA. Previous works by Vlasic et al. (2023) and Willis (2024) further confirm the applicability of quantum-inspired feature selection in resource constrained scenarios.

**Formatted:** Font: Lora

It must be noted that QIFS was applied to 100 samples alone, which were computationally selected, from the v1.0-mini dataset. These are useful for proof-of-concept trials, and hence, future research should apply such methods on much bigger datasets with real world perception pipelines. Additionally, this paper harnessed only LiDAR-based spatial features (x, y, z, length, width, height), leaving out radar and camera data that would ordinarily go into a full-scale sensor-fusion pipeline; such a constraint somewhat limits the generalizability of the conclusions to real multi-sensor systems. On the other hand, hardware acceleration backends like true quantum processors or GPUs may very well prove the speed-ups touted by QIFS algorithms in larger-scale studies that are still to be undertaken.

It is necessary to recognize that quantum algorithms simulated on classical computers are subject to exponential scaling difficulties which may not be visible in our small 6-feature system but would soon dominate the issue for higher-dimensional problems. Classical simulation of quantum algorithms does not avail the benefits of qubit superposition and entanglement thus the resulting costs of computation increase to the order of the problem size exponentially. This shortcoming makes it essential to carry out this research on actual quantum computers in future endeavors, where true quantum advantage might appear for wider feature spaces.

| Method | Features Selected | Interpretation | Pros | Cons |
|---|---|---|---|---|
| **PCA** | 3 PCs | Linear combination of all features | High retention of variance | Low interpretability |
| **MI** | 4-5 features | 'x', 'y', 'length', 'width', 'height' | Fast, simple | Ignores interaction effects |
| **RFE (LogReg)** | 5 features | 'y', 'z', 'length', 'width', 'height' | Effective for linear models | Model-dependent, greedy |
| **QUBO + QAOA** | 3 features | 'x', 'y', 'width' | Optimized globally, interpretable | Higher setup cost |

Table 42. *Brief comparison overview of feature selection methods used in sensor fusion for autonomous vehicle perception, including interpretability, efficiency, and practical considerations.*

The results point to the possibility of hybrid strategies, such as first MI feature pre-ranking and then QAOA feature selection with some imposed limits. Feature selection methods for real-time autonomous systems should understandably be accurate in their results, should be efficient enough to run under time constraints, and furthermore be flexible to cater to any changes in sensor configuration and environmental context as they advance toward bigger autonomy and edge computing.

In contrast, while classical ways such as RFE and MI are still strong baselines, quantum-inspired algorithms in their infancy show evidence of surpassing them in very limited instances when there needs to be a balance between optimization complexity and interpretability. Hence, if quantum computing sees further development, there may also be another impetus to focus on quantum-inspired methods for autonomous car perception systems.

**Formatted:** Font: Lora

While classical solvers treat QUBO and QAOA, the methods are 'quantum-inspired' regarding their theoretical tie to quantum computing paradigms. Their performance here speaks of engineering gains much before the actual hardware is popularly available.

**Conclusion:**

This study explored whether quantum-inspired feature selection methods may improve sensor fusion in autonomous vehicle perception. The results suggest that QIFS generates two small yet informative feature subsets in the presence of practical constraints by putting into contrast quantum-inspired methods such as Quadratic Unconstrained Binary Optimization (QUBO) and the Quantum Approximate Optimization Algorithm (QAOA) with traditional methods like Principal Component Analysis (PCA), Mutual Information (MI), and Recursive Feature Elimination (RFE). Using x, y, width, QAOA offered a very efficient feature subset that utilized fewer features but did not discard important information, whereas traditional methods such as RFE and MI uncovered five important spatial features: x, y, length, width, and height.

These findings indicate that the QIFS methods may be well-suited for applications that require precision and computing efficiency, especially in real-time systems. This is on par with studies by ~~Vasicetal~~Vlasic et al. (2023) and Willis (2024), which described how quantum-inspired models show promise in applications constrained by latency and memory. The importance of interpretable selection in safety critical conditions is also justified by a symbolic implication of the chosen features (Nahata & Oth~~t~~man, 2023; Yeong et al., 2025).

**Formatted:** Font: Lora

Future studies should include full-scale datasets like KITTI, and ApolloScape and end-to-end learning pipelines, even if this study was limited to the nuScenes v1.0-mini dataset and independent selection pipelines. Moreover, the benefit of QAOA runtime could be confirmed through execution on actual quantum hardware or advanced simulators, which might deal with the exponential scaling limitations associated with classical simulation of quantum algorithms~~Additionally, the advantage of QAOA runtime could be validated by executing it on real quantum hardware or advanced simulators~~. Another promising solution would be using hybrid pipelines, such as MI for preselection and QUBO for refinement (Wang, 2022; Pham & Raahemi, 2025).

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black

Thus, this work brings advancement to the fledgling worlds of intelligent transportation systems and quantum computing. Ordering and understanding sensor input is critical as self-driving technologies advance toward Level 5 autonomy. Not only are the quantum-inspired technologies unique, but they are also useful, which makes them significant assets in the development of reliable, real-time autonomous perception.

**Acknowledgements:**[redacted]ld like to express my sincere gratitude to [redacted] for his invaluable guidance and unwavering support throughout this research project. I am extremely grateful for [redacted] for his help in coordinating key aspects of [redacted] and for providing insights and constructive feedback along the way. Additionally, I would like to mention [redacted] for her inspiring knowledge. I'd also like to acknowledge the expertise and hard work that has gone behind the creation of my cited sources and the developers of the nuScenes datasets for open-access data. Special thanks to my peers for their thoughtful feedback and encouragement.

**References:**


- Baek, H., Kim, D., & Kim, J. (2023). Fast quantum convolutional neural networks for low-complexity object detection in autonomous driving applications (arXiv:2401.01370). arXiv. https://doi.org/10.48550/arXiv.2401.01370
- Benedetti, M., Lloyd, E., Sack, S., & Fiorentini, M. (2019). Parameterized quantum circuits as machine learning models. Quantum Science and Technology, 4(4), 043001. https://doi.org/10.1088/2058-9565/ab4eb5
- Caesar, H., Bankiti, V., Lang, A. H., Vora, S., Liong, V. E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., & Beijbom, O. (2020). nuScenes: A multimodal dataset for autonomous driving. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 1618–11628. https://doi.org/10.1109/CVPR42600.2020.0116
- Chang, Y. I. (2025). A survey: Potential dimensionality reduction methods (arXiv:2502.11036). arXiv. https://doi.org/10.48550/arXiv.2502.11036
- Chen, C., Rosa, S., Miao, Y., Lu, C. X., Wu, W., Markham, A., & Trigoni, N. (2019). Selective sensor fusion for neural visual-inertial odometry. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 10534–10543. https://doi.org/10.1109/cvpr.2019.01079
- Elaziz, M. A., Ahmadein, M., Ataya, S., Alsaleh, N., Forestiero, A., & Elsheikh, A. H. (2022). A quantum-based chameleon swarm for feature selection. Mathematics, 10(19), 3606. https://doi.org/10.3390/math10193606
- Farhi, E., Goldstone, J., & Gutmann, S. (2014). A quantum approximate optimization algorithm (arXiv:1411.4028). arXiv. https://doi.org/10.48550/arXiv.1411.4028
- Grant, E., Wossnig, L., Ostaszewski, M., & Benedetti, M. (2019). An initialization strategy for addressing barren plateaus in parametrized quantum circuits. Quantum, 3, 214. https://doi.org/10.22331/q-2019-12-09-214


**Formatted:** Font: Lora

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black, , Ligatures: None

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** Hyperlink, Font: (Default) Lora, 11 pt,

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black, , Ligatures: None

**Field Code Changed**

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black, , Ligatures: None

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black, , Ligatures: None

**Formatted:** Hyperlink, Font: (Default) Lora, 11 pt,

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black, , Ligatures: None

**Field Code Changed**

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

- Guyon, I., Weston, J., Barnhill, S., & Vapnik, V. (2002). Gene selection for cancer classification using support vector machines. Machine Learning, 46(1–3), 389–422. https://doi.org/10.1023/A:1012487302797
- Hira, Z. M., & Gillies, D. F. (2015). A review of feature selection and feature extraction methods applied on microarray data. Advances in Bioinformatics, 2015, 1–13. https://doi.org/10.1155/2015/198363
- Kannamarlapudi, H., & Chintalapudi, S. (2025). Quantum artificial intelligence for secure autonomous vehicle navigation: An architectural proposal (arXiv:2506.16000). arXiv. https://doi.org/10.48550/arXiv.2506.16000
- Khaleghi, B., Khamis, A., Karray, F. O., & Razavi, S. N. (2013). Multisensor data fusion: A review of the state-of-the-art. Information Fusion, 14(1), 28–44. https://doi.org/10.1016/j.inffus.2011.08.001
- Khan, M. A.-Z., & Al-Karaki, J. (2025). QI-MPC: A hybrid quantum-inspired model predictive control for learning optimal policies (arXiv:2504.13041). arXiv. https://doi.org/10.48550/arXiv.2504.13041
- Kraskov, A., Stögbauer, H., & Grassberger, P. (2004). Estimating mutual information. Physical Review E, 69(6), 066138. https://doi.org/10.1103/PhysRevE.69.066138
- LeCun, Y., Kavukcuoglu, K., & Farabet, C. (2010). Convolutional networks and applications in vision. Proceedings of 2010 IEEE International Symposium on Circuits and Systems, 253–256. https://doi.org/10.1109/iscas.2010.5537907
- Machine Design. (2015). Saved by the sensor: Vehicle awareness in the self-driving age. Machine Design. https://www.machinedesign.com/mechanical-motion-systems/article/21836344/saved-by-the-sensor-vehicle-awareness-in-the-self-driving-age
- Nahata, D., & Othman, K. (2023). Exploring the challenges and opportunities of image processing and sensor fusion in autonomous vehicles: A comprehensive review. AIMS Electronics and Electrical Engineering, 7(4), 271–321. https://doi.org/10.3934/electreng.2023016
- Pham, T. H., & Raahemi, B. (2025). An adaptive feature-based quantum genetic algorithm for dimension reduction with applications in outlier detection. Algorithms, 18(3), 154. https://doi.org/10.3390/a18030154
- Preskill, J. (2018). Quantum computing in the NISQ era and beyond. Quantum, 2, 79. https://doi.org/10.22331/q-2018-08-06-79
- Rattan, A., Rudra Pal, A., & Gurusamy, M. (2025). Quantum computing for advanced driver assistance systems and autonomous vehicles: A review. IEEE Access, 13, 17554–17582. https://doi.org/10.1109/access.2025.3532958


**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** Hyperlink, Font: (Default) Lora, 11 pt,

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black, , Ligatures: None

**Field Code Changed**

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black, , Ligatures: None

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** Hyperlink, Font: (Default) Lora, 11 pt,

**Field Code Changed**

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black, , Ligatures: None

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

- Su, Y., Shi, X., & Song, T. (2025). Research on vehicle safety based on multi-sensor feature fusion for autonomous driving task. Computers, Materials and Continua, 83(3), 5831–5848. https://doi.org/10.32604/cmc.2025.064036
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need (arXiv:1706.03762). arXiv. https://doi.org/10.48550/arXiv.1706.03762
- Vlasic, A., Grant, H., & Certo, S. (2023). An advantage using feature selection with a quantum annealer. Quantum Information and Computation, 23(13–14), 1091–1104. https://doi.org/10.26421/QIC23.13-14-2
- Wang, Y. (2022). When quantum computation meets data science: Making data science quantum. Harvard Data Science Review. https://doi.org/10.1162/99608f92.ef5d8928
- Willis, J. M. (2024). QIXAI: A quantum-inspired framework for enhancing classical and quantum model transparency and understanding (arXiv:2410.16537). arXiv. https://doi.org/10.48550/arXiv.2410.16537
    - ~~Wang, Y. (2022). When Quantum Computation Meets Data Science: Making Data Science Quantum. Harvard Data Science Review. https://doi.org/10.1162/99608f92.ef5d8928~~
- ~~Pham, T. H., & Raahemi, B. (2025). An Adaptive Feature-Based Quantum Genetic Algorithm for Dimension Reduction with Applications in Outlier Detection. Algorithms, 18(3), 154. https://doi.org/10.3390/a18030154~~
- ~~Elaziz, M. A., Ahmadein, M., Ataya, S., Alsaleh, N., Forestiero, A., & Elsheikh, A. H. (2022). A Quantum-Based Chameleon Swarm for Feature Selection. Mathematics, 10(19), 3606. https://doi.org/10.3390/math10193606~~
- ~~Rattan, A., Rudra Pal, A., & Gurusamy, M. (2025). Quantum Computing for Advanced Driver Assistance Systems and Autonomous Vehicles: A Review. IEEE Access, 13, 17554–17582. https://doi.org/10.1109/access.2025.3532958~~
- ~~Yeong, D. J., Panduru, K., & Walsh, J. (2025). Exploring the Unseen: A Survey of Multi-Sensor Fusion and the Role of Explainable AI (XAI) in Autonomous Vehicles. Sensors, 25(3), 856. https://doi.org/10.3390/s25030856~~
- ~~Vlasic, A., Grant, H., & Certo, S. (2023). An Advantage Using Feature Selection with a Quantum Annealer. Quantum Information and Computation, 23(13 & 14), 1091–1104. https://doi.org/10.26421/QIC23.13-14-2~~
- ~~Chen, C., Rosa, S., Miao, Y., Lu, C. X., Wu, W., Markham, A., & Trigoni, N. (2019). Selective Sensor Fusion for Neural Visual-Inertial Odometry. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 10534–10543. https://doi.org/10.1109/cvpr.2019.01079~~
- ~~Yeong, D. J., Velasco-Hernandez, G., Barry, J., & Walsh, J. (2021). Sensor and Sensor Fusion Technology in Autonomous Vehicles: A Review. MDPI AG. https://doi.org/10.20944/preprints202102.0459.v1~~


**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** Hyperlink, Font: (Default) Lora, 11 pt, , Ligatures: None

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black, , Ligatures: None

**Field Code Changed**

**Formatted:** Font: (Default) Lora, 11 pt, Font color: Black, Ligatures: None

**Formatted:** Indent: Left: 0.5", No bullets or

**Formatted:** Bulleted + Level: 1 + Aligned at: 0.25" + Indent at: 0.5"

**Formatted:** Font: (Default) Lora

**Formatted:** List Paragraph, Bulleted + Level: 1 + Aligned at: 0.75" + Indent at: 1", Don't adjust space between Latin and Asian text, Don't adjust space between Asian text and numbers

**Formatted:** Font: (Default) +Body (Calibri), 12 pt, Font color: Auto

**Formatted:** List Paragraph, Left, No bullets or numbering, Font Alignment: Auto

- ~~LeCun, Y., Kavukcuoglu, K., & Farabet, C. (2010). Convolutional networks and applications in vision. Proceedings of 2010 IEEE International Symposium on Circuits and Systems, 253–256. https://doi.org/10.1109/iscas.2010.5537907~~
- ~~Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2023). Attention Is All You Need (No. arXiv:1706.03762). arXiv. https://doi.org/10.48550/arXiv.1706.03762~~
- ~~Baek, H., Kim, D., & Kim, J. (2023). Fast Quantum Convolutional Neural Networks for Low-Complexity Object Detection in Autonomous Driving Applications (No. arXiv:2401.01370). arXiv. https://doi.org/10.48550/arXiv.2401.01370~~
- ~~Hira, Z. M., & Gillies, D. F. (2015). A Review of Feature Selection and Feature Extraction Methods Applied on Microarray Data. Advances in Bioinformatics, 2015, 1–13. https://doi.org/10.1155/2015/198363~~
- ~~Su, Y., Shi, X., & Song, T. (2025). Research on Vehicle Safety Based on Multi-Sensor Feature Fusion for Autonomous Driving Task. Computers, Materials and Continua, 83(3), 5831–5848. https://doi.org/10.32604/cmc.2025.064036~~
- ~~Willis, J. M. (2024). QIXAI: A Quantum-Inspired Framework for Enhancing Classical and Quantum Model Transparency and Understanding (No. arXiv:2410.16537). arXiv. https://doi.org/10.48550/arXiv.2410.16537~~
- ~~Khan, M. A.-Z., & Al-Karaki, J. (2025). QI-MPC: A Hybrid Quantum-Inspired Model Predictive Control for Learning Optimal Policies (No. arXiv:2504.13041). arXiv. https://doi.org/10.48550/arXiv.2504.13041~~
- ~~Kannamarlapudi, H., & Chintalapudi, S. (2025). Quantum Artificial Intelligence for Secure Autonomous Vehicle Navigation: An Architectural Proposal (No. arXiv:2506.16000). arXiv. https://doi.org/10.48550/arXiv.2506.16000~~
- ~~Nahata, D., & Othman, K. (2023). Exploring the challenges and opportunities of image processing and sensor fusion in autonomous vehicles: A comprehensive review. AIMS Electronics and Electrical Engineering, 7(4), 271–321. https://doi.org/10.3934/electreng.2023016~~
- ~~Zhang, Y., Carballo, A., Yang, H., & Takeda, K. (2023). Perception and sensing for autonomous vehicles under adverse weather conditions: A survey. ISPRS Journal of Photogrammetry and Remote Sensing, 196, 146–177. https://doi.org/10.1016/j.isprsjprs.2022.12.021~~
- ~~Grant, E., Wossnig, L., Ostaszewski, M., & Benedetti, M. (2019). An initialization strategy for addressing barren plateaus in parametrized quantum circuits. Quantum, 3, 214. https://doi.org/10.22331/q-2019-12-09-214~~
- ~~Machine Design. (2015). Saved by the Sensor: Vehicle Awareness in the Self-Driving Age. Machine Design. https://www.machinedesign.com/mechanical-motion-systems/article/21836344/saved-by-the-sensor-vehicle-awareness-in-the-self-driving-age~~
- ~~Chang, Y. I. (2025). A Survey: Potential Dimensionality Reduction Methods (No. arXiv:2502.11036). arXiv. https://doi.org/10.48550/arXiv.2502.11036~~

- ~~Wold, S., Esbensen, K., & Geladi, P. (1987). Principal component analysis. Chemometrics and Intelligent Laboratory Systems, 2(1–3), 37–52. https://doi.org/10.1016/0169-7439(87)80084-9~~
- ~~Caesar, H., Bankiti, V., Lang, A. H., Vora, S., Liong, V. E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., & Beijbom, O. (2020). nuScenes: A Multimodal Dataset for Autonomous Driving. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 11618–11628. https://doi.org/10.1109/CVPR42600.2020.0116~~
- ~~Farhi, E., Goldstone, J., & Gutmann, S. (2014). A Quantum Approximate Optimization Algorithm (No. arXiv:1411.4028). arXiv. https://doi.org/10.48550/arXiv.1411.4028~~
- ~~Khaleghi, B., Khamis, A., Karray, F. O., & Razavi, S. N. (2013). Multisensor data fusion: A review of the state-of-the-art. Information Fusion, 14(1), 28-44. https://doi.org/10.1016/j.inffus.2011.08.001~~
- ~~Preskill, J. (2018). Quantum Computing in the NISQ era and beyond. Quantum, 2, 79. https://doi.org/10.22331/q-2018-08-06-79~~

**Formatted:** List Paragraph, Space After: 0 pt

**Formatted:** List Paragraph

## Response to Reviewer 1

**Manuscript:** "Quantum-Inspired Feature Selection Regarding Sensor Fusion in Autonomous Vehicle Perception"

**Date:** December 9, 2025

---

Dear Reviewer 1,

I sincerely thank you for your thorough review and constructive feedback. Your input has greatly improved this manuscript. I go over each of your queries and recommendations in detail below, along with the modifications made to the updated manuscript.

-----------------------------------------------------------------------------------

**Comment 1:** To improve readability, please introduce each term only at its first mention and use the abbreviation thereafter. In particular: quantum-inspired feature selection (QIFS); Principal Component Analysis (PCA); Quadratic Unconstrained Binary Optimization (QUBO); Quantum Approximate Optimization Algorithm (QAOA). After the first definition, please use only the abbreviations throughout.

**Response:** I am grateful to the reviewer for pointing out this crucial issue related to consistency. I have evaluated the manuscript for the use of abbreviations and corrected it wherever necessary. I tried to be consistent all along the document, but I admit that some inconsistencies were left behind. Now, I have made it a practice that every technical term (QIFS, PCA, QUBO, QAOA, RFE, MI) is introduced with a full definition at its first appearance followed by its abbreviation in parentheses and that the abbreviation is only used in all subsequent occurrences throughout the manuscript.

**Changes Made:**

-Abstract: First mentions include full terms with abbreviations, now

-Introduction: Consistent use of abbreviations after first defined

-Throughout manuscript: Corrected instances where full terms were repeated unnecessarily

-Deleted entire sentences in the methodology and results section as they contained full term and abbreviation repetition, which could be best understood by context and readable when references were made

-Reviewed usage of abbreviations in accordance with their definitions

-----------------------------------------------------------------------------------

**Comment 2:** Please convert Figure 2 to a standard three-line table and revise Table 1 and Table 2 into the same three-line format. Also, place each table's title (caption) above the table, aligned with common academic style guides.

**Response:** Thank you for your advice on formatting. Figure 2 has been converted into a three-line table that is properly formatted and Tables 1 and 2 have been reformatted to standard academic style with horizontal lines only at the top, between headers and data, and at the bottom. All captions of tables have been moved above their corresponding tables according to the standard academic conventions.

**Changes Made:**

-Figure 2: Converted to Table format in Word itself with three-line style (top, header separator, bottom lines only)

-Table 1: Reformatted to three-line style with caption rightly placed

-Table 2: Reformatted to three-line style with caption rightly placed

-All vertical lines removed from tables per standard academic formatting

--------------------------------------------------------------------------------

**Comment 3:** The content currently describing dataset selection in lines 94–101 would fit better in the Materials & Methods section. Please move it there and add a brief example (e.g., a small snippet or schematic) to illustrate what a typical sample looks like, so readers can quickly grasp the data characteristics.

**Response:** I agree that the details on the dataset selection are better suited to the Materials & Methods section. I have moved this content and enlarged it with an example figure that illustrates the typical data structure. Thank you to the reviewer for pointing out this ambiguity issue.

**Changes Made:**

-Move dataset description (lines 94-101) to Materials & Methods section

-Added Table 1 showing sample LiDAR data structure with actual normalized values from the nuScenes v1.0-mini dataset

-Included description of the 11 raw features present in the dataset

-Clarified which 6 engineered features were actually used in the experiments

--------------------------------------------------------------------------------

**Comment 4:** The paper frames the work as improving sensor fusion, yet only LiDAR features are used and camera/radar are not included (line 206-208). This might be confusing to readers. What does sensor fusion mean in the article? Is it the fusion of multiple modalities or just the fusion of multiple LiDAR sensors? Please be clear and add the corresponding description in the article.

**Responses:** I express my gratitude to the reviewer for spotting this confusion. I do not deny that the reference to “sensor fusion” could have interpreted incorrectly. In this study, I am solely using the spatial features coming from one LiDAR sensor, and I am not doing multi-modal sensor fusion (combining LiDAR, camera, and radar). I have introduced a special clarification part, and I have also changed the way of presenting the idea in the paper.

**Changes Made:**

-New subsection added: “Clarification on Sensor Fusion Scope” in Materials & Methods (after Dataset Selection section)

-A clear statement: "The term ‘sensor fusion’ in this research denotes the combination of the spatial features derived from a single LiDAR sensor; it does not denote the combination of several sensor modalities, such as LiDAR, camera, and radar."

-In Introduction, statement added acknowledging that: "Although multi-sensor modalities (LiDAR, radar, and cameras) are available in nuScenes, this study was focused solely on LiDAR-derived spatial features (XYZ coordinates and object dimensions)"

-Conclusion now unmistakably points to multi-modal fusion as an important future research direction

---

**Comment 5:** The authors say intensity is used when forming feature vectors (line 213-218), but later define the engineered feature set as only spatial and size terms. Clarify exactly which features enter each method.

**Responses:** I have made it clear that intensity values were not take into account in the feature selection experiments. The 6 engineered spatial and dimensional features (x, y, z, length, width, height) were the only ones used.

**Changes Made:**

-Materials & Methods section: Introduced explicitly clear sentence: "The raw LiDAR point cloud data still come with intensity levels. The designed feature set used in this work, however, relies solely on spatial coordinates (x, y, z) and object dimensions in length, width, and height. The derived intensity level from the raw point cloud data has never been part of the feature vectors for any of the feature selection methods."

-Dataset Selection section was made clearer that 6 engineered features were taken from the ‘translation’ and ‘size’ fields of sample_annotation.json

-Indication of 6 engineered features, not all 11 raw features, being used for experiments was added.

---

**Comment 6:** In line 373, there is also a phrase “underwater human class”. What does it mean?

**Responses:** I apologize for this overlooked mistake. The phrase was a typographical error due to the autocorrect feature. The phrase I intended to use was "underrepresented object classes." This has been amended in the entire manuscript.

**Changes Made:**

-Line 373 (Discussion section): "underwater human class" was replaced with "underrepresented object classes"

-Now the sentence is: "...emphasizing the importance of underrepresented object classes at both geometric and semantic levels"

---------------------------------------------------------------------------------

**Comment 7:** The authors acknowledge the small sample and lack of error bars. Runtime and memory comparisons are also not reported. Please add cross-validated results, error bars, and a clear runtime/memory table for each method on CPU and, if possible, GPU.

**Responses:** I appreciate the important suggestion. I have now run 5 separate implementations of each of the four feature selection methods and reported the run time measurements as mean ± standard deviation. I have included an extensive runtime comparison in Table 2 with error bars.

**Changes Made:**

-Table 2: Now includes runtime measurements with error bars (mean ± std) for all methods:

- PCA: 0.044 ± 0.002 s
- RFE: 168.5 ± 4.1 s
- MI: 23.0 ± 0.4 s
- QUBO + QAOA: 1.51 ± 0.94 s

-Added note to Table 2: "All methods averaged over 5 independent runs"

-Materials & Methods: Added description of measurement methodology: "All measurements performed on 100-sample nuScenes v1.0-mini subset in Google Colab with T4 GPU runtime (CPU execution)."

-Results section: Added Computational Performance Analysis subsection reporting these measurements

-Note: The use/mention of GPU acceleration was not relevant since the methods we applied (that is, scikit-learn implementations, Qiskit simulator) did not support GPU computation at all. This limitation was pointed out in the write-up.

---------------------------------------------------------------------------------

**Comment 8:** The text says PCA is applied, then RFE with logistic regression "after PCA" (line 297), yet RFE is later described as operating on six original features. This creates confusion about the exact pipeline. Also, the QIFS run constrains the selection to three features without a stated reason. Please provide a single pipeline diagram, state which features are passed to each step, justify the choice of subset size, and report downstream model performance for each selected subset.

**Responses:** Apologies for the misunderstanding caused by the pipeline description. I have made it clear now that the PCA, RFE, MI, and QUBO + QAOA methods were running parallel (not sequentially), each one working independently on the same 6 engineered features. A pipeline diagram (Figure 2) has been included, and I have justified the sizes of the feature subsets.

**Changes Made:**

-Added Figure 2: Complete processing pipeline diagram showing:

- o Stage 1: Data Extraction (raw LiDAR data)
- o Stage 2: Preprocessing (StandardScaler normalization)
- o Stage 3: Feature Selection (4 parallel methods)
- o Stage 4: Model Training (Logistic Regression, Random Forest, SVM)
- o (+) Stage 5: Evaluation (5-fold cross-validation)

-Materials & Methods, Processing Pipeline section: stated clearly: "Four methods were applied in parallel-" followed by bullet points for each method

-Rationale for feature subset sizes discussed in Materials & Methods:

- o PCA: "3 principal components (76.7% variance retained)" – variance-based criterion
- o RFE: "Selected 5 features" – standard iterative elimination
- o MI: "Identified top 5 features" – ranking-based selection
- o QUBO + QAOA: "Constrained to 3 features" – Incorporated explicit note: The QUBO + QAOA technique had a restriction of picking only 3 features, which was a way of demonstrating the method's capability to reach the highest level of dimensionality reduction together with classification performance…"

-Clarified that RFE worked on the 6 original engineered features, not on PCA components

-Results section: model performance comparison added (although not the main focus, I note that all methods yielded similar classification accuracy within the margin of 2%)

---------------------------------------------------------------------------------------

**Summary of Major Changes:**

1. Ensured consistent abbreviation usage throughout manuscript
2. Reformatted all tables to three-line style with captions above
3. Moved dataset description to Materials & Methods with example table

4. Added "Clarification on Sensor Fusion Scope" subsection
5. Clarified that intensity was NOT used; only 6 engineered features
6. Corrected "underwater human class" typo to "underrepresented object classes"
7. Added Table 2 with runtime measurements and error bars (5 runs each)
8. Added Figure 2 pipeline diagram showing parallel processing
9. Clarified feature subset size rationale
10. Removed confusion about PCA □RFE sequence (they're parallel)

---

I'm confident that these changes have greatly enhanced the clarity, precision, and reproducibility of our work. I trust that the manuscript is now up to the standards required for publication and I would like to thank you once again for your helpful comments.

Sincerely,

The Author

## Response to Reviewer 2

**Manuscript:** “Quantum-Inspired Feature Selection Regarding Sensor Fusion in Autonomous Vehicle Perception”

**Date:** December 11, 2025

---

Dear Reviewer 2,

I am grateful for your detailed and insightful review. Your remarks concerning the term “quantum-inspired” and the request for mathematical details have been especially useful. Below, I take up each of your major and minor comments, and I also indicate the changes that have been made to the manuscript in order to strengthen it.

--------------------------------------------------------------------------------

**Substantive Edits**

--------------------------------------------------------------------------------

**Comment 1:** I am not entirely sure why the paper employs QUBO and QAOA techniques over quantum machine learning methods (QML). QML should outperform virtually all other quantum methods for feature selection, much as it would in the classical case. Thus, why the specific focus on QUBO, for example?

**Response:** I appreciate the raising of this important question. I have added a new subsection named “Rationale for QUBO and QAOA selection” in the Materials & Methods section to explicitly address this choice. My reasoning includes:

1. The QUBO representations depict feature selection as a problem with two possible decisions (select/not select each feature) intrinsically.
2. QAOA serves as a connector between the classical and quantum worlds, it can be executed on classical simulators while being suitable for the upcoming quantum hardware (NISQ devices).
3. The constrained optimization is handled quite easily by these techniques (e.g., choosing precisely k features).
4. Depending on the small feature space (6 features), the use of QUBO/QAOA is easy to handle computationally on classical hardware, while the QML methods usually consume more quantum resources.
5. I do not deny that the development of quantum hardware might lead to more powerful QML techniques in the future.

**Changes Made:**

-The new subsection “Rationale for QUBO and QAOA Selection” (after Processing Pipeline section) was added to Materials & Methods

-The five rationale points as mentioned above were included

-A forecasting statement was added: “Ultimately, quantum hardware will be available for more extensive access; thus, the study would explore QML methods.”

-An acknowledgement note was added in the discussion section recognizing quantum machine learning (QML) as a key future direction

---

**Comment 2:** You should briefly describe what is in these real world data sets in this sentence: “nuScenes, which is what has been utilized here, KITTI, and ApolloScape…”

**Response:** I agree that this needs clarification. I have added in brief descriptions of each dataset in the Introduction.

**Changes Made:**

-Introduction: Expanded sentence to include: “This could be evaluated using real-world datasets such as nuScenes (multi-modal urban driving data from Boston and Singapore), which is what has been utilized here, KITTI (stereo camera and LiDAR data from Karlsruhe), and ApolloScape (large-scale urban scenes from Beijing) within standard sensor fusion settings.”

---

**Comment 3:** The Introduction should discuss how sensor fusion is performed in greater detail since this is a key focus of the paper.

**Response:** I have expanded the Introduction to provide more context on sensor fusion processes and challenges, including the computational burden of integrating high-dimensional multi-sensor data.

**Changes Made:**

-Introduction, paragraph 1: Added slightly more detailed explanation: “Through sensor fusion—the synthesis of data from different sensors to form a complete and united representation of the driving environment—the cars combine the advantages of each sensor type to create an exhaustive and unified representation of their driving environment…”

-Added discussion of overlapping sensor coverages (LiDAR, cameras, radars) and their different roles (environment mapping, collision detection, parking assistance)

-Figure 1 also clearly shows sensor coverage zones to visualize the fusion challenge

---

**Comment 4:** I'm a bit concerned about the terminology used in the paper and title for QUBO and QAOA methods… Quantum-inspired refers to algorithms adapted from fully quantum algorithms into classical forms… But, QUBO and QAOA are not quantum-inspired, they are truly quantum, meant to be run on quantum, not classical, hardware... Thus, if you are referring to these algorithms in your title, the title should be changed to just Quantum Algorithms.

**Response:** I deeply appreciate this critical observation about terminology. After careful consideration, we respectfully maintain the use of “quantum-inspired” for the following reasons:

1. **Established Usage in Literature:** The phrase "quantum-inspired" has become a common term across the machine learning and optimization literature, for it also applies to classical implementations of quantum algorithms. Some of the latest works (Wang, 2022; Pham & Raahemi, 2025; Elaziz et al., 2022) consider "quantum-inspired" to be the synonym for classical simulations of quantum algorithms.
2. **My Implementation:** I performed the QUBO and QAOA tasks with Qiskit’s classical simulator on common CPU hardware, as opposed to quantum hardware. The algorithms make use of quantum-inspired optimization techniques, however, they are implemented on classical systems without any superposition or entanglement.
3. **Bridge Concept:** The methods can be seen as a transition, they have been developed using quantum principles, although they can be currently implemented classically and will still be compatible with future quantum hardware.
4. **Clarity for Readers:** Using "quantum algorithms" in the title could cause a misconception among readers that the authors had conducted quantum hardware experiments, which was not the case.

But we still consider the reviewer's point about precision valid and thus we have added in lot of clarifications across the manuscript.

**Changes Made:**

-Introduction: Added clarification: “These methods are said to be quantum-inspired because they are quantum-powered optimization paradigms, essentially meaning they do not require quantum hardware (Benedetti et al., 2019; Farhi et al., 2014).”

-Materials & Methods: Added explicit note: “An additional note of importance is that since quantum-inspired methods were simulated by available classical computers using Qiskit, these do not directly use real quantum computing resources; rather, they are designed from paradigms inspired by quantum computation.”

-Discussion: Added statement: “While classical solvers treat QUBO and QAOA, the methods are ‘quantum-inspired’ regarding their theoretical tie to quantum computing paradigms. Their performance here speaks of engineering gains much before the actual hardware is popularly available.”

-Throughout: Consistently used “quantum-inspired” with appropriate qualifiers about classical implementation.

--------------------------------------------------------------------------------

**Comment 5:** Your discussion would be improved if you expanded upon “It consists of accurate object annotations which are accessible synchronously with the recording from LiDAR, radar, and camera sensors (Caeser et al., 2020).” to state what the different variables are (column headings) and what objects were being imaged in the dataset.

**Response:** I agree and have added substantial detail about the dataset structure, variables, and object classes.

**Changes Made:**

-Materials & Methods, Dataset Selection: Added comprehensive description: “The dataset consists of annotated 3D bounding boxes for a variety of different object classes, such as cars, people, bikes and traffic cones, which were recorded in the different driving situations found in Boston and Singapore.”

-Added Table 1 showing the 6 engineered features used in experiments, with sample normalized values. Also added description of object classes (cars, people, bike, traffic cons) and annotation fields.

--------------------------------------------------------------------------------

**Comment 6:** How many total features were used as inputs? 6? This should be clarified since the following sentence makes it sound like 6, which is a very low-dimensional space already. “Although multi-sensor modalities (LiDAR, radar, and cameras) are available in nuScenes, this study was focused solely on LiDAR-derived spatial features (XYZ coordinates and object dimensions).”

**Response:** I have clarified throughout that exactly 6 engineered features were used. I acknowledge this is a low dimensional space and have added justification for this choice as a proof-of-concept study.

**Changes Made:**

-Materials & Methods: Stated explicitly: “While radar and camera streams were also available in nuScenes, these flows were not used in the study. Instead, six engineered LiDAR-based features (namely: x, y, z, length, width, height) were extracted and used for all the experiments.”

-Added rationale: “This was a strategic choice in order to emphasize purely geometric and dimensional features, which are most closely related to object detection and classification.”

-Introduction: Added acknowledgement: “Considering the dataset size and a reliance mainly upon LiDAR-based 3D spatial features, the study must be taken as an exploratory step assessing the real-time suitability of QIFS for autonomous driving applications.”

-Discussion: Acknowledged limitation: “this paper harnessed only LiDAR-based spatial features (x, y, z, length, width, height), leaving out radar and camera data that would ordinarily go into a full-scale sensor-fusion pipeline”

---

**Comment 7:** The paper would be improved by providing sufficient background for the uninitiated reader to understand your key QUBO and QAOA algorithms. Providing equations is standard. It is also useful to describe what the other algorithms to which you compare do. In my field, again, we’d do this with equations.

**Response:** I have substantially expanded the mathematical descriptions of all four feature selection methods in the Materials & Methods section. While full equation level formulations were not incorporated for the methods, the revised text provides sufficient technical details on the fundamental optimization principles and algorithmic structures to support reader understanding.

**Changes Made:**

-Materials & Methods: Added mathematical descriptions for all four methods:

- o PCA: Eigenvalue decomposition, variance retention criterion
- o RFE: Iterative backward elimination, greedy optimization
- o MI: k-nearest neighbor density elimination (k=3), independence assumption
- o QUBO + QAOA: Binary variables, objective function maximizing MI scores with cardinality constraint (k=3), penalty coefficient λ= 1000, Qiskit statevector simulator with p=1 layer, COBYLA optimization

---

**Comment 8:** I don’t fully understand this statement: “Due to resource limitations, we were not able to conduct experimental runtimes or efficiency measurements on all techniques under comparison (e.g., CPU vs GPU). Therefore, the results should be used in the context of feature selection quality rather than speed performance.” If you ran all of the algorithms presented, you must know wall-times?

**Response:** I apologize for the confusion caused by this outdated statement. I have since conducted comprehensive runtime measurements and removed the statement about resource limitations.

**Changes Made:**

-Removed the confusion statement about resource limitations

-Added Table 2: Computational Performance Analysis with complete runtime data (mean ± std over 5 runs):

- PCA: 0.044 ± 0.002 s
- RFE: 168.5 ± 4.1 s
- MI: 23.0 ± 0.4 s
- QUBO + QAOA: 1.51 ± 0.94 s

-Added Results section: "Computational Performance Analysis" describing these measurements

-Noted execution environment: Google Colab, T4 GPU runtime, CPU execution, Python 3.10

---

**Comment 9:** If the QUBO results are the focus of the paper, there should be figures depicting those results, but there are none in the paper right now. You can consider including a figure that shows convergence or which final features were selected.

**Response:** I have added Figure 4 showing the Mutual Information scores for all features, which formed the basis of the QUBO optimization. I also added Table 3 clearly showing which features were selected by each method.

**Changes Made:**

-Added Figure 4: Bar chart of MI scores for all 6 features

-Enhanced Table 2 to clearly show selected features for each method (PCA: 3 components, RFE: x/y/length/width/height, MI: y/x/width/height/length, QAOA: x/y/width)

-Results: Added explicit statement identifying QAOA's selected features

---

**Comment 10:** As stated above, no mathematical description of how the QUBO or QAOA algorithms were performed was included. This is necessary for understanding what your algorithms actually did (and whether one would expect any quantum advantage if truly quantum-informed).

**Response:** I have added comprehensive mathematical and algorithmic descriptions in the Materials & Methods section as described in response to Comment 7 above.

**Changes Made:**

-Kindly see detailed response to Comment 7 above, which addresses this concern.

---

**Comment 11:** I don't think any results were presented for QAOA, so discussions of it should likely be deleted.

**Response:** I respectfully note that QAOA results were presented, though perhaps not prominently enough. I have enhanced the visibility of QAOA-specific results throughout the updated manuscript.

**Changes Made:**

-Table 2: Clearly shows "QUBO + QAOA" method with runtime (1.51 ± 0.94 s) and selected features (x, y, width)

-Results section: Dedicated paragraph on QAOA: "Lastly, Quantum-Inspired Feature Selection (QIFS) was achieved by employing a Quadratic Unconstrained Binary Optimization (QUBO) model and the Quantum Approximate Optimization Algorithm (QAOA). In the QUBO formula, mutual information scores are transformed into negative weights to be minimized… QAOA was configured with COBYLA as the classical optimizer… The optimizations revealed that, within the constraints applied, x, y, and width were the three most informative features."

-Discussion section: Analysis comparing QAOA's compact selection (3 features) versus classical methods (5 features)

-Conclusion: Highlighted QAOA's contribution: "Using x, y, width, QAOA offered a very efficient feature subset…"

---

**Comment 12:** I agree with this statement! "Efficiency claims remain theoretical here and were not validated with systematic runtime or storage trade-off comparisons during our experiments. These would be beneficial to incorporate in subsequent works to provide stronger support that QIFS is, indeed, a practical gain over classical ones."

**Response:** I thank the reviewer for this supportive comment. Following this suggestion, I have now conducted systematic runtime measurements and added them to the revised manuscript.

**Changes Made:**

-Added Table 2 with empirical runtime measurements for all methods (5 runs, with mean ± standard deviation)

-Results section now includes quantitative evidence: PCA (0.044 ± 0.002 s), QAOA (1.51 ± 0.94 s), MI (23.0 ± 0.4 s), RFE (168.5 ± 4.1 s)

-Discussion now references measured performance instead of theoretical claims

-Acknowledged that larger-scale validation and storage trade-offs remain necessary future work

---

**Comment 13:** "emphasizing the importance of underwater human class at both geometric and semantic levels." What does the underwater human class refer to?

**Response:** I sincerely apologize for this atypical typographical error. This should read as "underrepresented object classes". I have corrected this.

**Changes Made:**

-Discussion section: Changed "underwater human class" to "underrepresented object classes"
-Sentence now reads: "...emphasizing the importance of underrepresented object classes at both geometric and semantic levels"

---

**Comment 14:** I would like to see more evidence of how the QUBO converged on features and with what weights. Whether 3 or 5 features are selected often depends heavily on settings and convergence. Thus, it is useful to see how QUBO converged to 3 features, but no details are provided.

**Response:** I have added detailed description of the QUBO formulation and constraint enforcement mechanism. The selection of exactly 3 features only was enforced through a hard constraint (penalty $\lambda$=1000), not through convergence behavior.

**Changes Made:**

-Materials & Methods: Detailed QUBO description showing objective function maximizes MI scores while penalizing constraint violations, with k=3 features chosen for maximum dimensionality reduction

-Specified penalty coefficient $\lambda$=1000 for strong constraint enforcement, resulting in exactly 3 features selected

-Added design choice explanation: "restriction of picking only 3 features, which was a way of demonstrating the method's capability to reach the highest level of dimensionality reduction together with classification performance, thus maximally utilizing the computational power for the real-time based autonomous vehicle applications."

---

**Comment 15:** A key ingredient, which you do mention, is run-time. Quantum algorithms run on classical hardware are typically exponentially more expensive than classical algorithms, which dwarfs their advantage. This is not seen on small systems (because you can't see an exponential scaling with few data points), but one should technically take scaling into consideration when finaling comparing algorithms.

**Response:** I fully agree with critical observation. I have added discussion of scalability considerations and the exponential scaling limitation of quantum algorithms on classical hardware.

**Changes Made:**

-Discussion: Added paragraph on scalability: "quantum algorithms simulated on classical computers are subject to exponential scaling difficulties which may not be visible in our small 6-feature system but would soon dominate the issue for higher-dimensional problems."

-Conclusion: "QAOA runtime could be confirmed through execution on actual quantum hardware or advanced simulators, which might deal with the exponential scaling limitations associated with classical simulation of quantum algorithms"

--------------------------------------------------------------------------------

**Minor Edits**

--------------------------------------------------------------------------------

**Comment 1:** "Juxtaposing them against traditional filter methods such as mutual information and Principal Component Analysis" is not a sentence and thus needs to be rephrased.

**Response:** Corrected. This piece has been integrated into a complete sentence.

**Changes Made:**

-Abstract: Rephrased to grammatically complete sentence: "Comparing them against traditional filter methods such as Mutual Information (MI) and Principal Component Analysis (PCA), then, the main conjecture is…"

--------------------------------------------------------------------------------

**Comment 2:** "Autonomous driving (also known as level 5 cars)" should be autonomous driving (also known as level 5) cars."

**Response:** Corrected as suggested.

**Changes Made:**

-Introduction: Changed to "autonomous driving (also known as level 5) cars"

---

**Comment 3:** Should be “Figure 1 offers” not “Figure 1 offered.”

**Response:** Corrected to present tense as suggested.

**Changes Made:**

-Introduction: Changed “Figure 1 offered” to “Figure 1 offers”

---

**Summary of Major Changes:**

1. Added "Rationale for QUBO and QAOA Selection" section explaining choice over QML
2. Added dataset descriptions (nuScenes, KITTI, ApolloScape) in Introduction
3. Expanded sensor fusion explanation in Introduction
4. Addressed "quantum-inspired" terminology with extensive clarification
5. Added comprehensive dataset variable description and Table 1
6. Clarified 6 engineered features used throughout
7. Added mathematical descriptions of all four feature selection methods
8. Removed confusing resource limitation statement; added Table 2 with runtimes
9. Added Figure 4 showing MI scores, enhanced feature selection visibility
10. Enhanced QAOA results presentation throughout manuscript
11. Added detailed QUBO formulation and constraint enforcement description
12. Added scalability discussion acknowledging exponential scaling on classical hardware
13. Fixed all three minor grammatical issues

---

I belief that the extensive revisions have considerably increased the manuscript’s technical rigor, mathematical depth, and clarity of the quantum-inspired nature of our approach. I am particularly thankful to the reviewer for the highlighting of the importance of precision in use of terms and in the mathematical description. I trust that the revised manuscript is now up to your criteria and I would like to express my appreciation for your meticulous and positive critique.

Sincerely,

The Author

## Response to Reviewer 3

**Manuscript:** "Quantum-Inspired Feature Selection Regarding Sensor Fusion in Autonomous Vehicle Perception"

**Date:** December 12, 2025

---

Dear Reviewer 3,

I sincerely thank you for your critical and beneficial assessments. Your remarks on the orderly presentation, numerical support, and citation style have been essential to the fortification of the paper. I will now respond to your comments indvidually and indicate the corresponding alterations in the updated version.

-----------------------------------------------------------------------------------

**Comment 1:** Unclear logical flow in the introduction. It is recommended to restructure the section as follows: start with the motivation, introduce the current limitations, identify the research gaps, present the proposed method.

**Response:** I appreciate this structural guidance and have completely reorganized the Introduction to follow the best possible logical progression.

**Changes Made:**

-Introduction restructured with clear flow: (1) Motivation, (2) Current Considerations, (3) Traditional methods & gaps, (4) Proposed quantum-inspired approach, (5) Research question & contribution

-Added clearer transitions between sections

-Reorganized paragraphs to follow logical progression

-----------------------------------------------------------------------------------

**Comment 2:** Why was the nuScenes v1.0-mini subset chosen instead of the full dataset or KITTI/ApolloScape? Does the small sample size (100 samples) affect the statistical robustness of the conclusions?

**Response:** I thank the reviewer for this necessary methodological question. I have added explicit rationale for the choice of dataset and acknowledged limitations clearly regarding statistical robustness.

**Changes Made:**

-Added rationale in Dataset Selection: "This subset was designed so that the representative samples from different driving conditions were maintained and thus, it would be possible to conduct rapid prototyping and preliminary testing"

-Introduction: Added "exploratory step" acknowledgment

-Discussion: Added upfront assessment of sample size limitations and noise susceptibility

-Conclusion: Identified future work with larger datasets (KITTI, ApolloScape)

---------------------------------------------------------------------------

**Comment 3:** The abstract claims that QIFS "significantly improves computational efficiency" (Page 1, Lines 20–22), but no quantitative evidence is provided.

**Response:** The unsubstantiated claim regarding efficiency has been eliminated from the Abstract and it has been replaced by accurate, evidence-based assertions which are evidenced by my runtime measurements.

**Changes Made:**

-Abstract: Removed unsubstantiated "significantly improves computational efficiency" claim

-Abstract: Replaced with: "selected only 3 interpretable features ('x', 'y', 'width') with negligible loss in accuracy (≤ 2%) compared to baselines, while maintaining computational efficiency and preserving feature interpretability"

-Added Table 2 with quantitaive runtime evidence (PCA: 0.044 s, QAOA: 1.51 s, MI: 23.0 s, RFE: 168.5 s)

-Results: Added "Computational Performance Analysis" subsection

-Discussion: Evidence-based statement: "resulting demonstration tests the capacity of quantum-inspired algorithms to provide competitively compact and high-quality feature subsets..."

---------------------------------------------------------------------------

**Comment 4:** The QIFS method selected features x, y, and width. Why does the algorithm prefer these features?

**Response:** I have added in detailed explanation of why QAOA selected these specific features, correlating them to both the Mutual Information scores and the physical importance for autonomous vehicle perception.

**Changes Made:**

-Results: Added explanation linking to MI scores

-Discussion: Added physical interpretation: "The x and y values give exact 2D locations of the objects in the LiDAR image, which is a prerequisite for algorithms of path planning, collision detection, and obstacle avoidance."

-Figure 4 caption: Enhanced to show why QAOA selected these features

---

**Comment 5:** It is recommended that the author summarize the workflow using a flowchart to help readers understand it more easily.

**Response:** That is a great suggestion. I have built a detailed flowchart (Figure 2) that depicts the entire 5-stage processing pipeline from raw data to the evaluation step.

**Changes Made:**

-Added Figure 2: Five stage processing pipeline flowchart showing Data Extraction □ Preprocessing □ Feature Selection (4 parallel methods) □ Model Training □ Evaluation

-Added “Processing Pipeline and Feature Selection Workflow” subsection in Materials & Methods with step-by-step description

-Each pipeline stage explicitly referenced in text

---

**Comment 6:** The reference list needs uniform formatting. Several references are not presented in the same style as the others.

**Response:** I sincerely apologize for the inconsistent reference formatting. I have completely reformatted the entire reference list with APA 7$^{th}$ edition standards and maintained consistency throughout the formatting.

**Changes Made:**

-All 29 references reformatted to uniform APA 7th edition
-Alphabetically sorted A-Z by first author's last name
-Fixed all titles to sentence case (only first word and proper nouns capitalized)
-Removed "No." from all 8 arXiv papers (e.g., "No. arXiv:1411.4028" □"arXiv:1411.4028")
-Fixed Yeong et al. (2021): Changed "MDPI AG" □"[Preprint]"
-Added missing Benedetti et al. (2019) reference
-Fixed Vaswani et al. year: 2023 □ 2017
-Fixed Vlasic et al.: "23(13 & 14)" □ "23(13–14)"

---

**Summary of Major Changes:**

1. Completely restructured Introduction with clear logical flow:
   Motivation □ Limitations □ Gaps □ Proposed Method □ Research Question

2. Added explicit rationale for nuScenes v1.0-mini choice
   Acknowledged statistical robustness limitations
   Identified future work with larger datasets

3. Removed unsubstantiated efficiency claims from Abstract
   Added Table 2 with quantitative runtime evidence
   Replaced claims with evidence-based statements

4. Added detailed explanation of why QAOA selected x, y, width:
   Relationship to MI scores
   Physical significance for AV perception
   Enhanced Figure 4 caption

5. Created comprehensive Figure 2: 5-stage workflow flowchart
   Added "Processing Pipeline" subsection in Methods
   Step-by-step description of entire pipeline

6. Completely reformatted all 29 references to APA 7th edition:
   Alphabetically sorted
   Consistent sentence case titles
   Fixed arXiv formatting (removed "No.")
   Added missing Benedetti et al. reference
   Fixed Vaswani et al. year (2023 □2017)
   Uniform punctuation and style throughout

---

I believe that the extensive modifications have greatly elevated the manuscript in terms of clarity, logical flow, quantitative rigor, and professional presentation. The rewritten Introduction, new flowchart, quantitative evidence, and consistent reference formatting not only cover your points but also pull up the general quality of the paper by a large margin. Thank you to your detailed review and helpful suggestions, and I strongly hope the revised manuscript is meeting the publication standards now.

Sincerely,

The Author

The authors have done an excellent job addressing my feedback, particularly by clarifying their methodology and adding helpful diagrams that make the research much easier to follow. While it is evident that this is one of the author’s first academic articles, and I still have minor reservations regarding certain definitions, such as the scope of "sensor fusion", the revisions show significant progress. Given the extensive effort invested in this version, I am happy to **accept the manuscript without further revision**.

# Quantum-Inspired Feature Selection Regarding Sensor Fusion in Autonomous Vehicle Perception Review #2

Overview: The author has put substantial work into addressing the reviewer comments. The author has substantially expanded many details, including in the Methods and Results Sections, per our request. This has greatly improved the readability and precision of the paper. I am happy to accept it with minor revisions at this point.

Below, I note some continued confusion that the author has about quantum algorithms. QAOA is a quantum algorithm, not a quantum-inspired algorithm, that only shows advantage on quantum computers. People have formulated QUBO problems on classical computers since about the 1930s and have solved them with a variety of sometimes very clever classical methods. Solving QUBO problems on classical computers should thus show no advantage relative to current classical methods. One can certainly run these algorithms on classical computers, but this confers no advantage; we typically just do this for benchmarking before running on quantum hardware. I am willing to let this go because the authors are likely not quantum (or statistical) theorists who know this history, but some of the statements in the paper would be considered odd by those who are experts.

**Substantive Comments:**

1. I still disagree with your interpretation of quantum-inspired algorithms; in fact, your interpretation is incorrect. For instance, a quick look-up of the definition of quantum-inspired states that "Quantum-inspired algorithms are classical algorithms that borrow ideas, mathematical structures, or intuitions from quantum computing—but run entirely on classical hardware and do not require qubits or quantum devices." Your algorithms aren't classical. Your algorithms are quantum algorithms. You have just run them on classical hardware. We all run quantum algorithms on classical hardware, including UCC and VQE algorithms, when possible due to scaling limitations. That doesn't make those algorithms quantum-inspired. I am not going to hang on this point, but your use of Quantum-Inspired just is not correct. A quantum-inspired algorithm would be one like quantum random walks, which came from quantum algorithms, but has been translated into the classical sphere; most such algorithms come from the quantum information literature.
2. This discussion isn't quite right: "Quadratic Unconstrained Binary Optimization (QUBO) and the Quantum Approximate Optimization Algorithm (QAOA) evolved as classical solvers built around quantum computation paradigms." QAOA did not evolve as a classical solver; it was always assumed to be a quantum solver. Let's read the first sentence of the original Farhi paper on QAOA you cited as proof: "We introduce a **quantum** algorithm that produces approximate solutions for combinatorial optimization problems." Again, it is correct that you can run these algorithms on classical hardware (quantum simulators), but they have no advantage when you do that. QAOA was explicitly developed to realize advantages on quantum machines, when they emerge as reliable.
3. "These methods run on classical systems today." These systems can run on classical systems today. They have also been used repeatedly on quantum hardware over the past 15 years, including on quantum annealers. As a result, this sentence is not accurate. Similarly, "Moreover, these methods might become directly compatible with quantum hardware as it matures, making these methods a quite forward-looking option for large-dimensional problems like AV sensor fusion." is not correct. If you are arguing about whether they have yet provided advantage, that's one question, but your statement is that they are not compatible with quantum hardware, which is not true. They have frequently been run on quantum hardware for over a decade. See this tutorial, which underlies many published demonstrations: [https://docs.dwavequantum.com/en/latest/quantum_research/qubo_ising.html](https://docs.dwavequantum.com/en/latest/quantum_research/qubo_ising.html).
4.

**Smaller Edits:**

1. Abstract: “This could be evaluated using real-world datasets such as nuScenes (multimodal urban driving data from Boston and Singapore), which is what has been utilized here, KITTI (stereo camera and LiDAR data from Karlsruhe), and ApolloScape (large-scale urban scenes from Beijing) within standard sensor fusion settings.” I think it would be more straightforward and concise to say “This was evaluated using real-world datasets…” and get rid of the “which is what has been utilized here.” It isn’t clear from the wording whether you mean that KITTI and ApolloScape weren’t used here. The current wording could be improved.
2. In a few places, logistic was misspelled.

After reviewing the revision, my decision is:

The authors have addressed most of the reviewers' comments. I recommend this paper for publication.